


AFRL-SN-RS-TR-1999-135

Final Technical Report

June 1999



NCTI SOFTWARE PROTOTYPE & TECHNOLOGY DEMONSTRATION

Calspan - UB Research Center

Joseph A. O'Sullivan

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

19990726 042

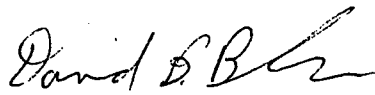
**AIR FORCE RESEARCH LABORATORY
SENSORS DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

DTIC QUALITY INSPECTED 4

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-SN-RS-TR-1999-135 has been reviewed and is approved for publication.

APPROVED:



DAVID B. BUNKER
Project Engineer

FOR THE DIRECTOR:



ROBERT G. POLCE, Acting Chief
Rome Operations Office
Sensors Directorate

If your address has changed or if you wish to be removed from the Air Force Research Laboratory Rome Research Site mailing list, or if the addressee is no longer employed by your organization, please notify AFRL/SNRT, 26 Electronic Parkway, Rome, NY 13441-4514.525. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE June 1999	3. REPORT TYPE AND DATES COVERED Final Feb 96 - Feb 97		
4. TITLE AND SUBTITLE NCTI SOFTWARE PROTOTYPE & TECHNOLOGY DEMONSTRATION		5. FUNDING NUMBERS C - F30602-93-D-0075, Task 26 PE - 63789F PR - 4072 TA - 08 WU - PN		
6. AUTHOR(S) Joseph A. O'Sullivan				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Calspan - UB Research Center 4455 Genesee Street Buffalo NY 14225		8. PERFORMING ORGANIZATION REPORT NUMBER N/A		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/SNRT (formerly Rome Laboratory) 26 Electronic Parkway Rome NY 13441-4514		10. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-SN-RS-TR-1999-135		
11. SUPPLEMENTARY NOTES Air Force Research Laboratory Project Engineer: David B. Bunker/SNRT/(315) 330-2345				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution unlimited.		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) The final report for Noncooperative Target Identification (NCTI) Software Prototype and Technology Demonstration discusses the efforts to improve on previously developed NCTI algorithms to allow them to be demonstrated using real data. This included modifying the code to accept real data provided by Rome Laboratory (RL), comparing the output of Xpatch to real data provided for a DC-10, modifying the code to run on many types of output data from Xpatch, extending the code to perform recognition based on sequences of High Resolution Radar (HRR) data, incorporating many more data types, and demonstrating through extensive simulations the ability to distinguish between targets using Xpatch generated data sets. These algorithms are not ready to exploit in a real world system but show much promise. Some additional work has also been done to gain on this goal. Among these are an improved model for the sensor to decrease sensitivity to orientation errors, comparing different models for the radar data, using simulated data at a variety of bands, and starting to model polarized data				
14. SUBJECT TERMS Noncooperative Target Identification, Algorithms, High Resolution Radar		15. NUMBER OF PAGES 170		
		16. PRICE CODE		
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Final Report to CUBRC: NCTI Software Prototype and Technology Demonstration

Joseph A. O'Sullivan
Department of Electrical Engineering
Washington University, Campus Box 1127
St. Louis, MO 63130

1 Introduction

Washington University has worked on a contract for \$ 34,250 for CUBRC, to support efforts on the Expert Science Program Task Order PR-A-6-1753, "NCTI Software Prototype and Technology Demonstration." This report has two parts. The first part is a brief summary of the work. The second is an appendix and is a complete copy of the doctoral thesis of Steven P. Jacobs. The work described in this doctoral thesis was supported in part by the CUBRC funds. Thus, this thesis describes not only work completed for CUBRC, but describes in detail a more comprehensive effort, extending the results done for CUBRC in significant ways.

In previous efforts, Washington University in St. Louis (WU) worked with Rome Laboratory for over three years on the project "Signal Processing Hierarchy for Radar and Communications," contract F30602-92-C-004. WU developed novel algorithms for hostile target identification. The approach addresses some limitations of current hostile target identification techniques.

In the proposal to CUBRC, the following items were identified as being necessary before the code developed at Washington University could be ready for running on real data:

- modify the code to accept real data inputs including requisite preprocessing
- perform a verification of the ability of Xpatch to predict real data
- modify the code to run on Xpatch data
- extend the code to perform recognition based on a sequence of range profiles

- incorporate more target types
- demonstrate the ability of the code to distinguish between these target types for Xpatch data sets

Washington University and Jim Llinas of the University of Buffalo have made progress in all of these areas. These are detailed below and in the appendix, and include:

- we have modified the code to accept real data provided by Rome Laboratory
- we have compared the output of Xpatch to real data provided for a DC-10
- we have modified the code extensively to run on many types of output data from Xpatch
- we have extended our code to perform recognition based on sequences of HRR data
- we have incorporated many more target types
- we have demonstrated through extensive simulations the ability to distinguish between targets using Xpatch generated data sets

In addition, we have made progress on other fronts including:

- we have improved the model for the sensor to decrease sensitivity to orientation errors
- we have compared two different models for the radar data, demonstrating their relative benefits
- we have extended our simulations to ground targets as well as airborne targets
- we have acquired and have used extensively a comprehensive database of simulated range profiles of four ground targets, the URISD
- we have used simulated data at a variety of bands
- we have started to model polarized data
- we have found an independent source for real data to be used in the future

We also met in St. Louis with personnel from Rome Laboratory.

These different areas of progress are outlined below and in the appendix.

2 Rome Laboratory Visit

On June 28, 1996, Vincent Vannicola, David Bunker, and Bill Baldygo of Rome Laboratory visited Washington University in St. Louis. The meeting included discussion of the project originally funded by Rome Laboratory that started in December 1991 and continued through July 1995. The final report for that project was discussed in detail including the interactions between Rome Laboratory and Washington University as part of that effort. Washington University has been involved in two efforts through CUBRC, including this project. Both were discussed.

For this project, the objectives of the project were discussed, including how they fit into the overall goals of research funded by Rome Laboratory. The areas in which Washington University, working with Jim Llinas, planned to make progress were outlined—these areas are listed in the introduction and discussed in more detail below.

3 Model Development for Improved NCTI Simulations

The algorithms developed by WU under previous Rome Laboratory sponsorship are designed to perform recognition of radar targets using high resolution radar (HRR) data and tracking data. The algorithms perform joint target tracking and recognition because the two problems of tracking and recognition, especially using HRR data, cannot be decoupled. There is a natural synergy between the tracking and recognition through their use of orientation information. As part of the recognition, the target orientation is estimated. The orientation estimates are fed to the tracker thereby improving tracking performance. If the tracking is improved, then the kinematic models for the target yield improved estimates of the target sequence of orientations, and hence improved prior information on the next orientation of the object. This improved prior then improves the orientation estimate at the next step. Since orientations are used in the recognition process, recognition performance increases.

The approaches developed under previous Rome Laboratory support and extended as part of this effort with support from CUBRC use Bayesian methods. There is a prior on the sequence of orientations that a target assumes in any given scene derived from a dynamical model for the target. The dynamical model is based on the standard Newtonian dynamics of a target in motion. In its most general form, there is a six degree of freedom model, yielding twelve differential equations for the motion. These twelve differential equations have six different inputs that summarize the inputs of including thrust, aerodynamic drag, aileron deflections, etc. The parameters within the model include effects derived from the size and shape of the aircraft including wing area, moments of inertia, mass, etc. Thus, in the extreme case, each aircraft may have its own

model for motion. In practice, aircraft are separated into classes such that each class has its own aerodynamic model. The inputs to the equations are modeled by random processes. These induce likelihood functions for the state process in the model, yielding priors on the orientations and positions of the aircraft.

The general approach described in the last paragraph may be applied to derive a prior on the sequence of orientations and positions using any appropriate dynamical model. It should be mentioned that in the standard tracking literature, much more simplified dynamical models are used. They are typically linear and have at most six states (positions and velocities). One of the differences that we promote vigorously (and has been proposed by others who are also looking at ways to merge high resolution data and low resolution tracking data) is the explicit modeling of the coupling between the orientation and the tracking model. Note that if one senses a change in orientation prior to a change in the velocity vector, one can anticipate the change in the velocity vector that will result. Thus, highly maneuverable targets may require orientation measurements and incorporation of orientations in the dynamics in order to achieve acceptable tracking performance. In our case, the measurements of orientation are through the HRR data which are very sensitive to target pose.

In the Bayesian approach, the prior is combined with a sensor likelihood model conditioned on the target position and orientation. This prior includes effects due to additive noise and due to the randomness inherent in the target signature. One of the major breakthroughs that has been made as part of the doctoral work of Steve Jacobs, which has been supported in part by this effort, is the proposal of using a conditionally Gaussian model for HRR data. This has led to a model that is much more robust to target orientation. That is, the traditional difficulty with using a deterministic model (as promoted by those who rely extensively on high quality signature prediction code) is that the range profiles decorrelate over small changes in aspect angle. For example, at X-band frequencies, this could be over two hundredths of a degree. On the other hand, the Gaussian model is robust on the order of five degrees. That is, the Gaussian model may require a factor of over 160 fewer templates than the deterministic model. One might expect that with that many fewer parameters, the performance may suffer. This is not the case, however. This fact is discussed in detail in Steve Jacobs' doctoral thesis that makes up a large part of this report. In there, joint tracking and recognition simulations are shown that indicate better performance using the conditionally Gaussian model.

We anticipate that this breakthrough will be key to developing realistic systems to perform recognition using HRR data. The conditionally Gaussian model requires less storage, requires fewer computations per recognition step, and performs better in joint tracking and recognition than the deterministic model.

4 Improvement of NCTI Simulations: Xpatch Data

The NCTI simulations have been modified to accept a wide variety of Xpatch generated data. These have used a DC-10 airframe CAD model that we purchased, X-29 and F-15 CAD models delivered to us as part of Xpatch, a HMMWV CAD model that has been adapted from a different simulator, and data from the University Research Initiative Synthetic Dataset (URISD) described below for their M1 tank, T1 tank, School bus, and Firetruck (ground targets). The models have been run and used data for UHF wideband, L-band, S-band, X-band, and near millimeter wave band (35 GHz). We have used different polarizations and multiple polarizations.

The DC-10 model was used to simulate HRR range profiles for comparison with the actual range profiles collected from a DC-10 as it flew by Rome Laboratory. Such data collected by the Rome Laboratory radar have no ground truth associated with them. These targets of opportunity do not have instrumentation on board so that the operator of the radar would know to within small orientation windows the exact relative orientations of the targets and the radar. As a result, the orientations used to generate the simulated data could not match the actual orientations. Washington University personnel tried to guess the relative orientations, but had little success. The simulator would need to know to within a fraction of a degree the orientation of the target in order to get a valid comparison of simulated Xpatch data and real data.

In connection with this, Washington University personnel discussed with personnel at MIT, at an Wright laboratory and at several Army laboratories their understandings of the ability of Xpatch to generate data that match real data. It has been the experience of many of them that Xpatch data have the same general characteristics of real data, but do not match the real data exactly. That is, the relative orientation changes needed to decorrelate real data and Xpatch generated data are approximately the same. Also, the dominance of some scattering sites is approximately the same. Many of the details of the range profiles generated by Xpatch are different from real data, however.

It was decided, with consultations from Rome Laboratory personnel, that the simulations should continue to pursue the problem of recognition based on Xpatch range profiles. The results should give indications of the ability of a system to perform joint tracking and recognition on real data. The main difference between a real system and that used in the simulations would be that the real system would be trained using real data rather than Xpatch generated data.

The simulations were extended to run on all of the targets described in this section. Extensive simulations were run to show the ability of the algorithm to perform joint tracking and recognition.

4.1 URISD

Washington University worked with Jim Llinas of the University of Buffalo and with several contacts at Rome Laboratory to identify suitable sources of target CAD models and sources of real data. In addition, Washington University personnel made contacts to attempt to find other sources of data.

One of these contacts was with researchers at the Massachusetts Institute of Technology, in their Laboratory for Information and Decision Systems, and at Boston University. Alan Willsky and his co-workers at MIT and Clem Karl of Boston University had access to the University Research Initiative Synthetic Dataset (URISD). Rome Laboratory personnel worked with us and with Air Force Office of Scientific Research personnel to gain approval successfully for Washington University to acquire the URISD. This dataset consists of simulated HRR returns for four simulated ground targets. The dataset is described in detail in Steve Jacobs' thesis in the appendix, so it is just summarized here. The four synthetic targets are an M1 tank, a T1 tank, a school bus, and a firetruck. Each target was illuminated by UHF wideband, L-band, and X-band signals. For each illumination, there were four polarizations collected (vv, vh, hv, hh). This was done at each of three elevations. At each elevation, HRR range profiles were collected over the entire circle of azimuth (at a sampling rate increasing with increasing carrier frequency).

While the URISD is not ideal from the perspective of air-to-air encounters, it does constitute a comprehensive dataset that can be used for extensive simulations. It is the most complete dataset of HRR simulated data that we have found.

Described in the appendix are joint tracking and recognition simulations using the URISD.

4.2 Real Data

Washington University worked with Jim Llinas of the University of Buffalo to discuss with contacts at Rome Laboratory the possibilities of collecting new HRR data using the S-band sensor at Rome Laboratory and of identifying existing datasets that would be appropriate for use in the NCTI simulations developed at Washington University. The possibility of collecting new HRR data using the S-band sensor at Rome Laboratory was deemed to be too ambitious for this project. There were not enough funds devoted to accomplish this and in fact the S-band radar apparently was not available for such experiments.

Jim Llinas and contacts at Rome Laboratory working with us identified possible existing datasets at Rome Laboratory for possible use in an HRR simulation. In addition, Rome Laboratory had previously provided data to Washington University for use in this research. It was decided that the data already supplied to Washington University was of the same quality as the remaining data. An example of the type of computations performed on these data is

described on pages 66 and 67 of Steve Jacobs' thesis in the appendix.

In order for real data to be used in an extensive simulation effort, there must be HRR range profiles collected over a wide variety of target orientations for use in development of the library. These range profiles (let us call them the training data) must have extensive ground truth known about them, and provided to the developers of the algorithms. The ground truth must include precise measurements of the relative positions and orientations of the target and the sensor for each of the range profiles. In addition, the data must have tracks in them for use in training for joint tracking and recognition. A separate dataset should be available for testing any algorithms derived using the training data.

Washington University developed collaborations with researchers at MICOM as part of their efforts on a different project supported by the Army Research Office. As part of those collaborations, a dataset was identified that fulfilled the requirements described above. Unfortunately, this dataset (the ARTI-1 dataset) was not delivered to Washington University until after the completion of the CUBRC contract. Washington University researchers hope that this dataset will be used at some time in the future. There is one major limitation of using this dataset. The results will not be directly comparable to Xpatch data. The identities of the targets have been removed from the dataset because of clearance considerations. Thus, we will not be able to generate Xpatch data for direct comparison purposes. However, we hope to be able to demonstrate the algorithm working on real data, which would be a large step toward an actual implementation.

5 Conclusions

Washington University personnel, working with Jim Llinas of the University of Buffalo, and working on projects related to and extending this effort, have made significant progress on all of the areas identified in the original proposal. In addition, progress has been made in several related areas. Perhaps the most significant is the derivation of a new model for HRR range profile data that could change the way the algorithms for recognition from HRR data run. This will be pursued vigorously in future research.

In addition to pursuing the conditionally Gaussian model paradigm for target recognition, there are several other areas for potentially fruitful research pointed to by this project. These include running the simulations on the newly acquired real data set, determining the role of fundamental limits on the ability to estimate target orientation in the recognition problem, including polarized data in a more powerful way, and looking more seriously at multisensor fusion.

WASHINGTON UNIVERSITY
SEVER INSTITUTE OF TECHNOLOGY
DEPARTMENT OF ELECTRICAL ENGINEERING

AUTOMATIC TARGET RECOGNITION USING HIGH RESOLUTION RADAR
RANGE-PROFILES

by

Steven P. Jacobs, M.S.E.E., B.S.E.E.

Prepared under the direction of Professor Joseph A. O'Sullivan

A dissertation presented to the Sever Institute of
Washington University in partial fulfillment
of the requirements for the degree of

Doctor of Science

May, 1997

Saint Louis, Missouri

WASHINGTON UNIVERSITY
SEVER INSTITUTE OF TECHNOLOGY
DEPARTMENT OF ELECTRICAL ENGINEERING

ABSTRACT

AUTOMATIC TARGET RECOGNITION USING HIGH RESOLUTION RADAR
RANGE-PROFILES

by Steven P. Jacobs

ADVISOR: Professor Joseph A. O'Sullivan

May, 1997

Saint Louis, Missouri

Recognition of aircraft and ground targets from high resolution radar (HRR) range-profiles is a notoriously difficult problem, due in large part to the extreme variability in the data for small changes in target orientation. To achieve recognition in the presence of this variability, the problem is posed in the context of joint tracking and recognition of a target using a sequence of observed HRR range-profiles. The likelihood function for the scene configuration combines a dynamics-based prior on the sequence of target orientations with a likelihood for range-profiles given the target orientation. The recognition system performs joint Bayesian inference on the target type parameter and the sequence of target orientations at the observation times.

Successful recognition is critically dependent on an appropriate model for the HRR range-profiles. A deterministic model and a conditionally Gaussian model are

introduced, and the likelihood functions under each model for varying orientations and target types are compared. The comparison is extended to include both aircraft and ground targets, different radar frequency bands, and full polarimetric range-profiles. The results of these comparisons show better performance for the conditionally Gaussian model in terms of the potential for correct recognition when the orientation estimate has significant error.

Fundamental limits on the performance of estimators of target orientation are obtained for the two models in terms of a Hilbert-Schmidt lower bound on the expected errors. The bound is evaluated for each model using simulated data as a function of the intensity of the noise in the observations. This analysis provides a specific criterion for model selection for this problem.

Simulations are presented in which each of the sensor models is combined with a constant orientation rate model for the target dynamics to produce an algorithm for joint tracking and recognition using HRR data. Results from the simulations show the performance of the algorithm in the presence of additive noise, including the expected angular estimation error and the probability of correct recognition. The algorithm using the conditionally Gaussian model achieves superior performance while requiring significantly less memory.

Contents

List of Tables	xiv
List of Figures	xv
Acknowledgments	xvii
1 Introduction	1
1.1 Joint Detection, Tracking, and Recognition of Multiple Targets	2
1.2 Joint Tracking and Recognition Using HRR Data	5
1.3 Contributions	7
1.4 Outline of Thesis	8
2 Review of Literature	10
2.1 Conventional Methods for Radar-Based ATR	10
2.1.1 Historical Developments	11
2.1.2 Refinements and Extensions	12
2.1.3 Summary	15
2.2 Bayesian Inference for ATR	15
2.2.1 Sequential Hypothesis Testing	15
2.2.2 MAP Estimation for ATR	16
2.2.3 Hidden Markov Models	16
2.2.4 The Adaptive Gaussian Classifier	17
2.3 Orthogonal Subspace Projections	19
2.3.1 Eigenfaces and Fisherfaces	19
2.3.2 Eigenspace Methods for Pose Estimation	20
2.3.3 Classification of Commercial Aircraft	21
2.4 Joint Tracking and Recognition	22
2.4.1 Joint Position and Orientation Estimation	23

2.4.2	Joint Tracking and Recognition by Sequence Comparison Methods	23
2.5	Other Methods Using Radar Data	25
2.5.1	Recognition Using Correlation Filters	25
2.5.2	Recognition Using Wavelet Representations	26
2.5.3	Modulation by Propulsion Substructure	27
2.6	Radar Reflectivity Models	28
2.6.1	One-Dimensional Models	28
2.6.2	Three-Dimensional Models	29
2.6.3	Stochastic Models	31
2.7	Conclusions	33
3	High Resolution Radar Models	35
3.1	Basic Concepts in Radar Modeling	35
3.1.1	Point Targets	36
3.1.2	Range Resolution and Range-Spread Targets	38
3.1.3	Doubly-Spread Targets	40
3.1.4	The Range Profile	42
3.2	Modeling Range-Profiles	44
3.2.1	Deterministic Model	46
3.2.2	Conditionally Gaussian Model	54
3.3	Incorporation of Multiple Polarizations	74
3.3.1	Wave Polarization	74
3.3.2	Extension of Conditionally Gaussian Model	77
3.4	Stochastic Reflectivity Model for Correlated Range-Profiles	81
3.5	Conclusions	83
4	Orientation Estimation Using HRR Data	85
4.1	The Rotation Matrix	85
4.2	Orientation Invariance of HRR data	87
4.3	Estimation of Line-of-Sight Rotation Using Polarimetric Data	92
4.4	Performance Metrics for Orientation Estimation	96
4.5	Kinematic Priors for ATR	101
4.5.1	Recognition Using Kinematic Measurements	102
4.5.2	Kinematic Priors in Multisensor ATR	104
4.6	Kinematic Models for ATR Using HRR	105

4.7	Conclusions	107
5	Simulations	108
5.1	Joint Tracking of Position and Orientation for a Single Maneuvering Aircraft Target	108
5.2	Joint Orientation Estimation and Recognition for a Closing Aircraft Target	112
5.3	Joint Orientation Tracking and Recognition of Ground Targets	119
5.3.1	Observations	120
5.3.2	Libraries	120
5.3.3	Estimation Algorithm	121
5.3.4	Simulation Results	122
5.4	Conclusions	127
6	Conclusions and Future Directions	129
6.1	Conclusions	129
6.2	Future Directions	131
	References	134
	Vita	143

List of Tables

3.1	Frequency and azimuth parameters for the URISD.	49
-----	---	----

List of Figures

3.1	Illustration of high resolution radar.	39
3.2	Simulated L-band range-profiles for four ground vehicles.	50
3.3	Deterministic model log-likelihood function for a four-vehicle recognition test using L-band data.	51
3.4	Simulated X-band range-profiles for four ground vehicles.	52
3.5	Deterministic model log-likelihood function for a four-vehicle recognition test using X-band data.	53
3.6	Conditionally Gaussian model log-likelihood functions for four-vehicle recognition tests.	56
3.7	Conditionally Gaussian model log-likelihood functions for four-vehicle recognition tests.	58
3.8	Comparison of conditional Gaussian log-likelihood functions for noise variances of 0 and 0.01 square meters.	60
3.9	Comparison of conditional Gaussian log-likelihood functions for noise variances of 0 and 0.1 square meters.	61
3.10	Simulated X-band range-profiles for four ground vehicles near 90 degrees azimuth.	63
3.11	Log-likelihood functions for four-vehicle recognition tests near 90 degrees azimuth.	64
3.12	Real millimeter-wave range-profiles of a ground vehicle of unknown type.	65
3.13	Simulated millimeter-wave range profiles for two ground vehicles.	66
3.14	Log-likelihood functions for a two-vehicle recognition test using simulated millimeter-wave range-profiles.	67
3.15	Real S-band range-profiles for three commercial aircraft.	68
3.16	Model comparison using real HRR range-profiles.	69
3.17	Simulated S-band range-profiles for two aircraft models.	70

3.18	Deterministic model log-likelihood functions for a two-aircraft recognition test.	72
3.19	Conditionally Gaussian model log-likelihood functions for a two-aircraft recognition test.	73
3.20	Conditionally Gaussian model log-likelihood functions for four-vehicle recognition tests using polarimetric data and diagonal polarimetric covariance.	78
3.21	Conditionally Gaussian model log-likelihood functions for four-vehicle recognition tests using polarimetric data and non-diagonal polarimetric covariance.	80
4.1	Plots of $\sin 2\alpha$ and $\frac{1}{2} \sin \left(4\alpha + \frac{\pi}{6}\right)$	95
4.2	Hilbert-Schmidt bound on the expected squared error in azimuth estimation for the "M1" tank model viewed at 30 degrees azimuth. . . .	99
4.3	Hilbert-Schmidt bound on azimuth estimation error for the "M1" tank model viewed at 30 degrees azimuth.	101
5.1	Results of orientation estimation for simulation 1.	111
5.2	Illustration of local search algorithm.	115
5.3	A frame of graphics output from simulation 2.	116
5.4	Results of orientation estimation for simulation 2.	117
5.5	Results of joint orientation estimation and recognition for simulation 2. . . .	118
5.6	Data from a typical realization of simulation 3.	122
5.7	Estimation results from a typical realization of simulation 3 using L-band data.	123
5.8	Performance measures for simulation 3 using L-band data.	124
5.9	Performance measures for simulation 3 using X-band data.	126

Acknowledgments

I wish to thank my wife Julie Fiez for her unwavering love and encouragement. This thesis was completed in large measure because she believed in me even when I doubted myself. I also wish to give thanks to my parents Paul and Nancy Jacobs for their guidance and support throughout my life and the years of my graduate education.

I thank my advisor Professor Joseph O'Sullivan for providing me with the opportunity to work on this exciting and ambitious project. The research comprising this thesis would not have been possible without his guidance, expertise, and patience. He continually made himself available for my repeated questions, and pushed me to make this thesis the best it could be.

Professors Daniel Fuhrmann, Bijoy Ghosh, Michael Miller, and Donald Snyder deserve my thanks for their participation on my dissertation committee and for their critical analysis of this manuscript. I am indebted to Professor Snyder and Professor Miller for their guidance and numerous consultations during the preparation of my dissertation, and for allowing me to return to ESSRL for the completion of my doctoral program. I also want to acknowledge Professor Fuhrmann for directing my first graduate research project, and for his consultation on matters both professional and personal over the years.

I express my considerable gratitude to the students of ESSRL. I feel particularly fortunate to have worked with Anuj Srivastava, who taught me a great deal about a wide variety of subjects relating to object orientation and recognition. I would like to thank Mohammed Faisal, Anuj Srivastava, Matthew Cooper, and Zoltan Bekker for contributing to the various computer simulations that are part of this dissertation. Finally, my thanks to Don Porter, Mihir Sathe, Nick Cutaia, Dakshi Agrawal, Aaron Lanterman, and Paritosh Tyagi for their consultations and companionship.

This work was supported in part by Rome Laboratory F30602-92-C-0004, ARO DAAH04-95-1-0494, ONR N00014-95-1-0095, CUB 1756G, and CUB 1757A.

Steven P. Jacobs

Washington University in Saint Louis
December 1996

Chapter 1

Introduction

A problem of considerable interest within the field of estimation theory is the recognition of objects from observations made with remote sensors. This problem appears in a variety of application-specific contexts including automated manufacturing, autonomous vehicle navigation, and face recognition. In this thesis, we will be concerned with a particular variant of this problem known as automatic target recognition or ATR. In the ATR problem, the goal is to classify ground vehicles or aircraft as viewed through a variety of remote sensors. This has applications both in a military context, in which the goal is to determine the identity of potentially hostile targets without relying on visual contact, and in civilian air traffic control, in which the goal is to monitor the identity of all aircraft in the vicinity of a given airport. Many of the conventional systems employed for such problems rely on the targets to communicate their identities explicitly to the recognition system. In contrast, we will address the problem in which the targets are non-cooperative and do not communicate their identities.

The problem of automatic target recognition receives a wide variety of formal definitions in the literature, depending on the specific application being addressed. These problems may be divided into two broad groups. Most researchers investigating the ATR problem consider a context in which a single object or target is to be recognized. This constitutes the first group of ATR problems. The second group, in which multiple targets are to be recognized using the observation of multiple sensors, is discussed in the next section. In this first group, the object position is known to a high degree of accuracy, such that a high resolution sensor can be focused on the target, and a sequence of observations are recorded. Also, it is often assumed that the sensor has an unoccluded view of the target. The goal is to classify the target

as one of a finite set of possible target types, based on the observations collected. Loosely stated, the problem is "if I am staring at it, can I recognize it?" The sensor data are assumed to exhibit a predictable dependence on the target shape in order for recognition to be successful. Sensor data in the form of imagery or radar echoes are generally dependent on the orientation of the target relative to the sensor. Thus the problem of ATR in this context becomes one of providing a characterization of the data one expects to observe from a given target at varying orientations, such that when novel data from a target are obtained, the identity of the target may be determined. Clearly such systems will be dependent on an appropriate mathematical model for the data as determined by the physics of the sensor. The primary goal of this thesis is to investigate models for a high resolution radar (HRR) sensor, and to evaluate these models in terms of the ability to recognize targets using radar data.

1.1 Joint Detection, Tracking, and Recognition of Multiple Targets

In a military scenario, we must consider a considerably more general recognition problem than that discussed above. In this version of the ATR problem, multiple targets are encountered whose positions, orientations, and indeed presence or absence are unknown *a priori*. This constitutes the second broad group of ATR problems. A system for ATR to be deployed in such a scenario must be able to recognize all of the targets present in such a complicated scene. Researchers in the Center for Imaging Science have posed the problem in this general context, and have proposed a very ambitious solution. In order to account for all of the variability present in the scene, the proposed system is required to detect all targets as they enter or leave the scene, to track the positions and orientations of all targets as they move within the scene, and to recognize all targets. Multiple sensors are employed for this purpose, including low resolution sensors such as tracking radars that contribute primarily to the detection and tracking of targets, and high resolution sensors such as high resolution radars, forward-looking infrared or FLIR cameras, and optical imaging systems that contribute primarily to estimation of the target orientation and recognition. The goal of the system is to perform inference on the complete configuration of the scene, given all of the available sensor data.

A very complete exposition of the system proposed for ATR in the general context is given jointly in the overview by Miller [47] and in the dissertation of Srivastava [76]. The foundations for the scene representation and the resulting inference algorithms are derived from *pattern theory* as developed by Grenander [30]. The representation for the scene configuration is developed in the natural parameter space for all of the targets. An aircraft target will have, for any fixed time, a position that takes its value in \mathbb{R}^3 and an orientation that is represented by a 3×3 rotation matrix that takes its value in the group $\text{SO}(3)$. Also associated with the target is a discrete label a_1 taking its value in the discrete space \mathcal{A}_1 . As we will ultimately be interested in estimating the target positions and orientations at the times when they are observed by the various sensors, as well as recognition of the target, we form a track for each object as the concatenation of these positions and orientations, together with the target type label. Thus a track for an aircraft target observed at n_1 discrete times takes its value in the space

$$(\mathbb{R}^3 \times \text{SO}(3))^{n_1} \times \mathcal{A}_1. \quad (1.1)$$

A scene inhabited by m_1 aircraft, such that the summed length of all m_1 tracks is given by $n_1(m_1)$, is an element of

$$(\mathbb{R}^3 \times \text{SO}(3))^{n_1(m_1)} \times \mathcal{A}_1^{m_1}. \quad (1.2)$$

Ground vehicles present in the scene are assumed to move on flat ground, such that their position lies in \mathbb{R}^2 and their orientation is described by a 2×2 rotation matrix in $\text{SO}(2)$. Let the target type label for ground vehicles a_2 take its value in the set \mathcal{A}_2 . Extending the previous notation, a scene inhabited by m_1 aircraft and m_2 ground vehicles has tracks $x(m_1, m_2) \in \mathcal{X}(m_1, m_2)$, where

$$\mathcal{X}(m_1, m_2) = (\mathbb{R}^3 \times \text{SO}(3))^{n_1(m_1)} \times \mathcal{A}_1^{m_1} \times (\mathbb{R}^2 \times \text{SO}(2))^{n_2(m_2)} \times \mathcal{A}_2^{m_2} \quad (1.3)$$

Since the number of targets is unknown *a priori*, the complete parameter space for the scene configuration is

$$\mathcal{X} = \bigcup_{m_1=0}^{\infty} \bigcup_{m_2=0}^{\infty} \mathcal{X}(m_1, m_2). \quad (1.4)$$

By posing the ATR problem in this general context, we are required to perform inference on a parameter space that is very large and complex. However, this also allows us to constrain the possible estimates of the scene configuration by taking advantage of known relationships between the various parameters. The inference algorithm employed in this context empirically reconstructs the posterior density for the scene configuration conditioned on the observations through random sampling. Let y denote the complete history of all sensor measurements and let x denote the scene configuration to be inferred. Then the posterior density $p(x|y)$ may be decomposed through Bayes' rule as

$$p(x|y) \propto p(x)e^{L(y|x)}, \quad (1.5)$$

where $p(x)$ is the prior on the scene configuration x , $e^{L(y|x)}$ is the likelihood function for the data, and the constant of proportionality is independent of x . The prior $p(x)$ includes a complexity prior on the number of targets as well as a prior on the kinematic state of the target that is derived from an appropriate model for the target motion. In the case of aircraft targets, the Newtonian equations of motion produce a series of twelve differential equations relating the position and orientation of the target to the forces and torques acting on it. Given a sequence estimates of the position and orientation derived from previous observations, the equations of motion are used to induce a prior on the position and orientation of the target at the time of the next observation. The maximal use of prior information in this way allows us to construct an inference algorithm that visits candidate scene configurations of high posterior probability.

Multiple sensors are used to view the scene. A tracking sensor may be included; the resolution of this sensor is assumed to be larger than the target dimension, so that the data produced are independent of the target shape and orientation. For recognition, we have investigated the use of several high resolution sensors, and examined the improved performance when using multiple sensors to view the same target. The high resolution sensors under consideration include an optical imaging sensor such as a video camera, a FLIR camera, and a high resolution radar. Models for each of the sensors are developed and are used to derive likelihood functions for the observations conditioned on the parameters. A primary component of this thesis is the derivation of two candidate models for the high resolution radar sensor, and the comparison of the recognition and orientation estimation performance under these models.

The inference algorithm developed for ATR by Miller, et al. [31, 47] and Srivastava [76] seeks conditional mean estimates of the scene configuration under the

posterior density. The key to this approach is to construct a random sampling algorithm that follows jump-diffusion dynamics, such that empirical averages generated by the algorithm converge to the conditional mean under the posterior. The jump-diffusion algorithm consists of discrete jumps in the parameter space that are executed at random exponential times, and between these jumps the algorithm proceeds continuously according to a diffusion process. The jump moves are used both to estimate the naturally discrete parameters such as the number of targets and the target types, and also to allow the algorithm to make large changes in the estimates of the continuous parameters, when such changes are supported by the data. The diffusion process smoothly deforms a history of position and orientation estimates following a stochastic differential equation. Details concerning the construction of the jumps and the diffusion process can be found in [47, 76]. The key to the success of this approach is to choose the parameters of the diffusion process such that the posterior measure μ on the complete parameter space is also the unique stationary measure of the jump-diffusion process we construct, and the process so constructed has the ergodic property that the sample averages generated converge to the expectation under the posterior. A theorem establishing this important ergodic result is given in [76].

The approach as outlined above is the appropriate context in which to frame the ATR problem. The algorithm is computationally intensive and requires multiple sensors. However, it has several important advantages that make it attractive, including:

- multiple mobile targets are accommodated,
- the simultaneous solution of the tracking and recognition problems allows the estimation of target type to improve tracking and vice-versa, and
- fusion of data from multiple sensors occurs naturally within the Bayesian framework.

1.2 Joint Tracking and Recognition Using HRR Data

Work on development of an ATR system under this approach clearly requires work in several areas, only some of which are the subject of this thesis. For further information concerning closely related work conducted in the areas of sensor characterization,

kinematic priors derived from dynamical models for target motion, and the recognition problem in general, see [23, 50, 48, 47, 60, 62, 61, 77, 78, 76]. Specifically, this thesis addresses the modeling of high resolution radar data, optimal processing of the data for estimation of orientation and target type as derived from these models, and performance of estimators for jointly estimating target type and a sequence of target orientations from a sequence of observations. In order to investigate these issues, we have examined a simplified ATR problem in which the goal is to jointly estimate a sequence of orientations and the target type from a sequence of observed range-profiles. Addressing this problem allows us to focus on specific issues relating to the HRR data.

Consider a single target that has been detected and tracked, such that an HRR sensor can be properly aimed at it. The target is illuminated by a sequence of radar pulses, and the received echoes, known as range-profiles, are collected. Let $r_k(t)$ denote the k th range-profile, and θ_k the orientation of the target during the k th illumination. The collection of observations R_k and the sequence of orientations Θ_k up to the time of the k th measurement are given by

$$R_k = \{r_1(t), r_2(t), \dots, r_k(t)\} \quad \Theta_k = \{\theta_1, \theta_2, \dots, \theta_k\}. \quad (1.6)$$

Also, let $a \in \mathcal{A}$ be a discrete label indicating the target type.

The joint posterior density for the measurements conditioned on the parameters is decomposed through Bayes' rule into the product of a prior density on the parameters and the likelihood of the observations given the parameters,

$$p(\Theta_k, a | R_k) \propto p(\Theta_k, a) p(R_k | \Theta_k, a). \quad (1.7)$$

Because the orientations θ_k and the observations $r_k(t)$ occur sequentially, we may write

$$p(\Theta_k, a | R_k) \propto p(\theta_0, a) \prod_{\ell=1}^k p(\theta_\ell | \Theta_{\ell-1}, a) e^{L(r_\ell | R_{\ell-1}, \Theta_\ell, a)}, \quad (1.8)$$

In the cases of a deterministic model for $r_k(t)$ or conditionally independent observations, the posterior may be written as the product

$$p(\Theta_k, a | R_k) \propto p(\theta_0, a) \prod_{\ell=1}^k p(\theta_\ell | \Theta_{\ell-1}, a) e^{L(r_\ell | \theta_\ell, a)}, \quad (1.9)$$

where $p(\theta_0, a)$ is a prior on the initial orientation and the target type, $p(\theta_\ell | \Theta_{\ell-1}, a)$ is a prior on the ℓ th orientation, given all previous orientations and the target type, and $L(r_\ell | \theta_\ell, a)$ is the log-likelihood of the ℓ th observation, given the ℓ th orientation and the target type.

The prior $p(\theta_\ell | \Theta_{\ell-1}, a)$ is derived from an appropriate dynamical model for the target motion. Note that we cannot make use of the kinematic prior derived from the equations of motion as used in [47, 76], because we are assuming no availability of tracking information. In the simulations of Chapter 5, we will consider simplified dynamical models that allow us to place deterministic constraints on the sequence of estimates produced by our algorithm.

Of crucial importance to the performance of an estimator for this problem is the behavior of the data log-likelihood L for varying orientations and different target types, which in turn depends on the model chosen for HRR data. A primary factor making recognition from range-profiles difficult is the extreme variability in the observed data for small changes in target orientation. This represents a considerable modeling challenge as the model must accurately predict the variation, but the resulting recognition algorithm should be robust with respect to this variation. The choice of possible radar data models and the relative merits of two candidate models are discussed in Chapter 3.

1.3 Contributions

The original contributions of the research comprising this dissertation relate to the use of the high resolution radar sensor for orientation estimation and recognition. A deterministic model for HRR data, which is well-known and is presented for example in Van Trees [86] and Smith [73], is applied to simulated data, and the resulting likelihood function for varying orientations and target types is examined. Previous efforts have used deterministic models for ATR which assume that the target consists of a small number of scattering centers, which is not assumed here. A conditionally Gaussian model is applied directly to complex-valued range-profile data, and the likelihood function is analyzed in terms of the potential for recognition. This is a model for range-profiles that has not appeared previously in the literature, and that yields better performance for recognition of targets than any model we have seen. The conditional Gaussian model is extended to the case of fully polarimetric measurements, and the likelihood functions under two models for the polarimetric

covariance matrix are compared to the likelihood using a single linear polarization channel.

Both the deterministic and conditionally Gaussian models are analyzed to derive bounds on the performance of orientation estimation using HRR data. This provides specific criteria for model selection for this problem. We know of no other effort in which deterministic and stochastic models for HRR data are compared in this context. A class of orientations that produce equivalent range-profiles under circular polarization is defined, and equations that transform an arbitrary orientation described by three Euler angles to a pair of angles that uniquely identify the equivalence class are derived. The problem of resolving the ambiguity resulting from this equivalence using linear polarization is investigated, and a necessary condition for the maximum likelihood estimate is obtained in the form of a transcendental equation.

Two simulations are presented demonstrating the potential for joint tracking and recognition using a simple dynamical model for the target motion and data from only a HRR sensor. This is different from the methods investigated by Libby [45, 44] in which the maximum likelihood estimates of the orientations for two consecutive observations are computed individually, and dynamic programming is used to generate a sequence of orientation estimates. Other distinctions include that the method of constraining the sequence of orientation estimates produced by Libby's algorithm is not derived from a dynamical model for the target motion, and Libby assumes the availability of tracking data in the form of range, range rate and pointing angle measurements which we do not. Other investigators working on this problem assume knowledge of the target orientation in order to perform recognition without specifying how that knowledge is to be obtained.

1.4 Outline of Thesis

The remainder of this thesis is organized as follows. Chapter 2 presents a review of the open literature pertaining to automatic target recognition using high resolution radar, including an extensive compilation of various techniques employed to address this problem, and a discussion of various models for HRR data that have been applied. Chapter 3 gives an introduction to concepts in high resolution radar, derives two candidate models for HRR data, and goes on to conduct a series of tests of the recognition performance under each model by examining the likelihood functions using simulated data. Chapter 4 presents a characterization of local and global errors in orientation

estimates obtained using HRR data, describes possibilities for resolving some global errors, and investigates fundamental limits on the performance of orientation estimation and recognition using examples of simulated HRR data. Chapter 5 describes three simulations we have developed demonstrating the potential for joint tracking and recognition using HRR data. Finally, Chapter 6 identifies the major results that have been obtained in this thesis and gives a listing of a variety of challenging areas in which continued work is required for this problem.

Chapter 2

Review of Literature

This chapter contains a review of the open literature relating to the issues addressed in this thesis. We begin by examining conventional approaches to the problem of recognition using high resolution radar data. This is followed by several efforts in which recognition is posed as a Bayesian inference problem, other efforts in which orthogonal subspace projections are used for recognition. We discuss the work of Libby, who is the only other author we have found who investigates joint estimation of a sequence of orientations and target type from a sequence of observed range-profiles. The chapter is concluded by examining deterministic and stochastic models for high resolution radar data that have been proposed for use in recognition.

It should be pointed out that there are a number of other references in the literature concerning recognition of targets using range-profile data that are not cited here. All of the references that we have been able to obtain which provide the results of recognition of realistic targets using high resolution radar range-profiles are described in this chapter. Some of the references we obtained that do not meet the above criterion were omitted. Also note that there is a great deal of work on the problem of recognition of objects using synthetic aperture radar (SAR) imagery. No commentary on these efforts will be given, other than to say that one of the problems associated with using SAR images is that it may take an unacceptably long time to collect the data for each image.

2.1 Conventional Methods for Radar-Based ATR

Over the years, a wide variety of algorithms have been proposed for identifying aircraft from radar reflections. This section describes a number of such methods that have

been termed “conventional,” because they are largely derived from three historical developments in radar-based ATR:

- The demonstration by Copeland that the polarization of reflected radar waves is related to a target’s global shape.
- The assertion by Ksienski and others that target-specific information can be extracted from a target via multi-frequency radar reflections at wavelengths on the order of the overall target dimension.
- The assumption by Moffat and others that the impulse response of a radar target can be characterized by a small number of orientation independent natural resonances.

This section describes each of these developments and later extensions and published results in using them for target identification.

2.1.1 Historical Developments

Polarization Features

One of the earliest proposed methods for classification of objects based on their reflection of radar transmissions was that of Copeland [22]. The target is modeled as a pair of coupled antennae, a linear reflector and a sphere, and data are collected by rotating a linearly polarized receiving antenna about the line-of-sight and measuring the complex received voltage as a function of rotation angle. These data are processed to extract three parameters that are used for identification: the axis ratio of the polarization ellipse for each of the two hypothetical antennae and the relative tilt angle between these ellipses. Copeland goes on to show that these parameters take on significantly different values for various simple reflectors and postulates that comparison of the parameters measured for a given target to these known values will allow classification among the various reflector types. This result is important because it demonstrates that the polarization response of a radar target is dependent on its shape.

Multi-frequency RCS

Ksienski [41] proposes a target identification scheme based on multi-frequency measurements of the radar cross section. A set of n measurements of a radar target at

a particular aspect angle is represented by a point in n -dimensional complex space. Classification of a target is performed either by a linear discriminant, in which the data space is divided into class-specific regions by hyper-planes, or by a nearest neighbor technique, in which classification is achieved by comparing an observed range-profile to a library of signatures recorded for a discrete set of aspect angles. The nearest neighbor technique assumes some degree of knowledge of the aspect which resulted in the observation. Results are reported from classification studies using experimentally obtained data for simple shapes and scaled models of existing aircraft. The primary result is that the linear discriminant is suitable for pairwise classification of these objects, but when it is desired to identify an observed target as belonging to a single class out of many, the nearest neighbor approach is superior.

Natural Resonances

A third method of radar-based target identification given by Moffat [53, 52, 54] is based on the assumption that the impulse response of a radar target is dominated by the effects of a small number of low-frequency natural resonances. The resonance frequencies and their amplitudes, which are assumed independent of both the target orientation and the polarization of scattering data, are computed from thin wire models of the targets of interest using Prony's method and stored as a library. Identification is performed by comparing the ramp response of the target synthesized from multi-frequency scattering measurements to a ramp response calculated using the stored natural resonances. Results are presented which show successful classification of three aircraft models so long as scattering data are available at a sufficient number of harmonically related frequencies.

2.1.2 Refinements and Extensions

Time Domain Classification

Chen [17] extends the multi-frequency RCS method by considering comparison of measurements and stored library elements not only directly in the frequency domain but also in the time domain through preprocessing using the inverse discrete Fourier transform. Classification in the frequency domain uses a nearest neighbor decision rule. Time domain classification is performed by maximizing the correlation between observed and prestored data. In either case, the orientation that resulted in the observations is assumed known. Experimental results are shown for data collected on

models of five ships and four aircraft for each of three different target orientations. The time domain classification scheme is shown to be superior in terms of misclassification percentage unless the number of frequencies used is extremely low.

The Extinction Pulse Technique

Rothwell and Chen [66, 18, 65, 19] extend the work of Moffat. They conclude that if the target impulse response could in fact be characterized by a few natural resonances, then an "extinction pulse" could be synthesized such that the late-time response of the target to this pulse would ideally be zero. The late-time response is defined as the "free-ringing" response after the transmitted signal traverses the target. Illuminating a target with the extinction pulse of a different target reveals the discrepancy through significant energy in the late-time. The identification system as actually implemented stores a version of the extinction pulse for each object of interest and convolves each of these with the response of an unknown target to an arbitrary signal. As with Moffat's work, the assumption here is that the resonances are independent of target orientation.

Citing insufficient robustness of Prony's method with respect to noise and assumed model order, two new methods are provided for determining the resonances. The first involves direct minimization via Newton's method of the convolution of an observed response with the extinction pulses corresponding to different resonances. The second is a direct solution for the resonance frequencies using the method of moments.

Experimental results are given in which thin cylinders of differing lengths and four simple aircraft models are discriminated using this technique. The authors claim that other experiments in which aircraft models are illuminated at a variety of orientations demonstrate the aspect independence of the extinction pulse, but these results are not published. The primary drawbacks of this technique, as reported by Borden [12], are that the energy in the late-time response is small, so that high signal-to-noise ratios are required for successful recognition, and not all resonances are excited at all target orientations, so the technique is not truly independent of orientation.

Transient Polarization Response

Chamberlain, [16] draws from the work of both Copeland and Ksienski in his development of the transient polarization response (TPR) as a signal appropriate for target discrimination. Noting the connection between the polarization response of an object and its shape, Chamberlain sees polarization information as crucial to identification. In particular, the TPR is defined as the horizontally and vertically polarized responses to a circularly polarized transmitted signal. This response is assumed to be dominated by a small number of scattering centers, for which the polarization ellipses can be extracted directly from the TPR. Rather than observe the time domain response directly, the TPR is computed via the inverse fast Fourier transform from a series of multi-frequency RCS measurements for wavelengths in the range first proposed by Ksienski. The ellipticity, tilt angle and amplitude of each polarization ellipse are collected as an aspect-dependent feature vector for the target. Identification is performed by comparing observed and prestored feature vectors using four different distance metrics. Simulation results are presented for the discrimination of five scale models of commercial aircraft (Concorde, DC-10, 707, 727, 747) using the TPR when the target orientation is known. The authors claim that satisfactory rates of correct classification are achieved for all four distance metrics and for all signal-to-noise ratios above 0 dB, although this is not obvious from their published results.

Steadly [79] extends this concept by estimating the dominant scattering centers parametrically using Prony's method and backward linear prediction, rather than directly via the IFFT. This technique is applied to compact radar range measurements made on the same aircraft models used by Chamberlain. Success in estimating the polarization properties of scattering centers is demonstrated, but no attempt is made at direct comparison with Chamberlain's results nor at using the estimated parameters for identification.

Silverstein [71] also uses Prony's method to estimate polarization parameters for the target. Citing that this method is susceptible to production of false scattering centers, the identification algorithm allows for two or more scattering centers to be combined when comparing observed data to stored templates. The remaining scattering centers are associated with one another or with nil structures to find the best match between observation and template. A distance measure is defined which takes these effects into account. Simulation results are presented for the same five aircraft models used by Chamberlain and Steedly. Unlike Chamberlain's results, significant

difficulty is encountered in classifying one of the five targets for sufficiently high noise levels.

2.1.3 Summary

Two common elements of all of the above conventional methods are apparent:

- They are not based on Bayesian estimation via a statistical model for the received data.
- They do not view the target as a dynamically changing object whose orientation must be tracked in order for recognition to be achieved. As such, they do not propose joint or concurrent estimation of orientation and target type.

In the following three sections we examine a number of efforts that make advancements on one or both of these topics.

2.2 Bayesian Inference for ATR

In this section, we consider several references in which the recognition problem is cast as a problem of Bayesian inference from range-profile data. While the problem of recognition from radar data has a long history, we have found relatively few examples in which the problem is addressed in a Bayesian framework, and even fewer in which the resulting algorithms are tested using realistic target models. All such efforts we are aware of are listed here.

2.2.1 Sequential Hypothesis Testing

The first treatment of the Automatic Target Recognition problem that we have found in which radar measurements are modeled as a sample waveform from a random process is provided by Therrien [84]. He applies sequential hypothesis testing to the binary classification problem and shows that the resulting algorithm is in fact a two stage process of linear prediction of a future observation and classification based on comparison of the resulting residual sequences. While all of the essential elements of the classifier are described, no attempt is made toward implementation on real or simulated data.

Jouny [37] extends this idea to the M-ary case. Additionally, the target classes are each divided into subclasses corresponding to a sampling of the possible target

orientations in ten degree intervals. Jouny describes four suboptimal sequential tests whose complexity is lower than the optimal sequential test. The resulting classifiers are tested using four-frequency backscatter measurements of the five commercial aircraft models also used by Chamberlain, Steedly, and Silverstein. Results showing performance of the four algorithms as a function of noise level are given. While the use of subclasses corresponding to different target orientations implies that estimates of the target orientation are in fact generated by this procedure, no results are presented with respect to the accuracy of the resulting orientation estimates, or of the effect of the coarse sampling in orientation angle.

2.2.2 MAP Estimation for ATR

Smith [73] provides a Bayesian treatment of the ATR problem, in which the target is hypothesized to belong to one of several known classes, and radar data are modeled as random. Identification is performed by selecting the hypothesis which has the highest probability of being true, conditioned jointly on the received data and any prior information. For the case where the received data are the sum of a deterministic, aspect-dependent radar reflection and a sample waveform from a white Gaussian random process, the equation for the maximum a posteriori classifier is derived. Implementation of this equation requires knowledge of the deterministic reflection portion of the received data for each possible target; the observed data enters the estimator through an inner product with each of these deterministic signals. This in turn requires that the orientation of the target during observation be known or estimated. Smith goes on to discuss the issue of developing a library of radar reflections and a scattering center representation for the data. No implementation of the resulting classification algorithm is reported.

2.2.3 Hidden Markov Models

DeWitt [24] investigates the use of Hidden Markov Models in classification of targets from synthesized range-profiles. Prony model parameters are estimated from these data using backward linear prediction and the singular value decomposition. For each target of interest, a Hidden Markov Model is trained to range-profile features corresponding to a $10^\circ \times 10^\circ$ sector in aspect angle. Adjustment of HMM parameters in response to training data is performed via the Baum-Welch re-estimation procedure.

Classification is performed by computing the probability of each HMM having produced an observed range-profile feature vector, and selecting the target whose HMM probability is highest. Orientation estimates obtained from a tracking system are used to select the appropriate HMM for each target. The resulting system, once all of the HMM's have been sufficiently trained, is capable of classifying a target based on a single observed range-profile. This system is augmented with additional HMM's which are trained to characterize the temporal relationships between range-profiles. Simulation studies are performed to compare the classification performance of the various HMM configurations with each other and with a K-nearest neighbor algorithm. The single range-profile classifier is seen to outperform the KNN algorithm. DeWitt expects better performance from systems including inter-range-profile relationships, but the published experiments show the opposite to be the case. While the precise reason for this negative result is not clear, it is noteworthy that the attempt made to exploit the temporal relationships between observations are in no way based on a dynamical model for the object motion.

2.2.4 The Adaptive Gaussian Classifier

The Adaptive Gaussian Classifier or AGC is a system developed by Hughes Aircraft Corporation for recognition of non-cooperative targets from HRR range-profiles currently under study by the Air Force Wright Laboratory. Much of the literature concerning this algorithm is of a classified nature, thereby complicating the evaluation of this tool for comparison with other methods. However, two overviews of this algorithm are provided by Eisenbies [26] and Mitchell [51] that describe the essential components including the cost function. From these descriptions, it is learned that the term "Gaussian" refers to the computation of the Gaussian log-likelihood function for identification purposes, and that the term "Adaptive" refers to the energy normalization and range-alignment procedures applied to incoming range-profiles prior to classification.

The AGC algorithm consists of several consecutively-applied operations that may be grouped into two categories under the general headings of pre-processing and classification. The pre-processing steps are as follows, and are applied in the order listed:

1. computing the magnitude of the complex data in each range bin,
2. downsampling,

3. coarse range-alignment using the circular centroid,
4. normalization to unit signal energy, and
5. power transformation.

The power transformation step is defined by [26, 51]

$$d[i] = (c[i])^{0.4}, \quad (2.1)$$

where $c[i]$ is the return in range bin i after energy normalization and $d[i]$ is the same return after the power transformation. The rationale here is that the processed range-profiles are assumed to be distributed according to a Rician probability density function, and that the transformation of equation (2.1) “make[s] the underlying UHRR radar pdf more closely appear Gaussian,” [26] so as to validate the use of the Gaussian log-likelihood.

The classification of range-profiles is performed by comparison with templates developed for each target class. Each template consists of a mean value $\hat{\mu}[i]$ and a variance $\hat{\sigma}^2[i]$ of the returns in each range bin i that are computed from training data. Note that the variance computation implicitly assumes that returns from different range bins are uncorrelated. While Eisenbies does not directly comment on the need to characterize the range-profile dependence on target aspect angle, it is known from other correspondence [1] that templates are developed over $5^\circ \times 5^\circ$ patches in azimuth and elevation angles for each target. Also, the AGC does not estimate target orientation from the observed range-profiles, even though the available templates might allow for such estimation. It is assumed that orientation estimates are provided by an external system.

Once the templates have been developed, pre-processed range-profiles are classified by comparison with each template. This consists of two steps:

1. fine range-alignment by applying circular shifts to the received range-profile until the Euclidean distance to the mean template is minimized, and
2. computation of the Gaussian likelihood energy given by

$$\ln |\hat{\Sigma}| + (\underline{x} - \hat{\underline{\mu}})^T \hat{\Sigma}^{-1} (\underline{x} - \hat{\underline{\mu}}), \quad (2.2)$$

where \underline{x} is the received range-profile and $\hat{\Sigma}$ is the diagonal covariance matrix whose non-zero entries are $\hat{\sigma}^2[i]$.

The template that minimizes the likelihood energy is selected as the target type estimate for range-profile \underline{x} . Eisenbies goes on to test the performance of the AGC in a controlled experiment, with the result that the AGC yielded performance slightly inferior to that of k-nearest neighbor and neural network classifiers.

2.3 Orthogonal Subspace Projections

In a variety of inference problems, it is often desired to estimate a small number of parameters from observations that take their value in a space of high dimension. In some cases, the data vectors for all possible values of the parameter lie in a low-dimensional subspace of the original data space. In such cases, it can be very advantageous to determine a projection from the data space to the low-dimensional subspace, and solve the inference problem using the projected data. In this section we consider efforts toward applying such techniques to the recognition problem. Significant progress has been made toward the recognition of faces in images using these techniques, so we begin by examining these efforts.

2.3.1 Eigenfaces and Fisherfaces

Recent efforts toward systems for recognition of faces in images include work by Belhumeur and colleagues [7, 8, 32] and Hallinan [33], in which a low-dimensional representation is sought for the set of images of a face or other object for varying illumination conditions. The rationale for this work is developed in [8], where it is demonstrated that the set of images of a face for all possible light source directions and intensities forms a 3-D linear subspace of the image space. An orthogonal basis for the linear subspace is found by computing the eigenbasis for a selected set of images under differing illuminations. The resulting basis vectors are themselves images and are referred to as eigenfaces. The principal eigenfaces can then be used to construct an image of the same face under arbitrary lighting conditions [8, 33] or, given a novel image of the same face, determine the lighting conditions under which it was produced [33].

The eigenface methods in [8, 33, 32] consider variation in lighting only; it is assumed that the images used have been controlled for other factors such as facial expression, translation, scale, and pose. By contrast, in [7] it is assumed that variation in facial expression produces a small deviation in the resulting images from

the linear subspace, such that linear projection into an orthogonal basis can be used for recognition of faces that is invariant both to lighting and expression. A single basis is computed for the set of face images over several individuals, and the image of an unknown face to be recognized is projected along the basis vectors. Recognition is performed by computing the distance in the projected feature space from the unknown face to each of the subspaces for the different individuals. The nearest subspace is chosen as the identity of the unknown face.

Given this method, it is demonstrated that the eigenfaces computed in [8, 32, 33] are not the best basis for recognition. Because a single basis is computed across all individuals, the principle components of the eigenbasis will contain the maximal variation in all parameters, including the lighting and expression to which the recognition system should be invariant. Belhumeur [7] proposes the use of a basis computed using Fisher's Linear Discriminant for recognition. This basis is chosen so as to maximize the ratio of the variation between individuals to the variation for each individual in the projected feature space. The resulting basis vectors are termed Fisherfaces. Results are presented that demonstrate superior performance for the Fisherface method over that of the eigenface method for lighting and expression invariant recognition.

It is important to note that the recognition system described above makes no attempt to deal with variation in pose. In [32], a system is described for tracking the region in an image containing a face when the face undergoes translation and rotation in the image plane, but none of these efforts consider extending orthogonal subspace methods to the case of general pose variation. The next subsection describes two recognition methods that attempt to account for such variation.

2.3.2 Eigenspace Methods for Pose Estimation

Murase and Nayar [55] consider the problem of joint recognition and pose estimation of simple objects from images. Images are obtained from several objects for a discrete sampling of the possible poses; each image is normalized in scale and brightness. The resulting images are used to compute a single "universal" [55] eigenbasis for the images of all objects at all poses. In the absence of highly specular objects, the authors reason that the object image will change smoothly with varying pose, and therefore that the projections of these images into the eigenspace will change smoothly as well. Thus, the set of images for an object at all possible orientations,

when projected into the eigenspace, will form a smoothly varying manifold. If the pose is described by a single orientation angle such as azimuth, then the manifold is a closed curve. The set of images for each object form a separate closed curve in the universal eigenspace. Recognition of a novel image is performed by projecting into the eigenspace and locating the object whose manifold is nearest the projected novel image. The location of the point on the selected manifold nearest the projected novel image determines the pose of the object in the image. Alternatively, once an object has been recognized, an individual eigenbasis is computed using the images from that object only, and pose estimation is performed in the object-specific eigenspace. Experiments are conducted for four-class and twenty-class recognition, with very good results in terms of recognition rate and pose error, given that a sufficient number of images are used to construct the bases.

Pentland, et al. [63] take a different approach to pose estimation. Rather than compute a universal eigenbasis for all objects at all poses or an individual eigenbasis for each object at all poses, the authors propose computing an eigenbasis across all objects at each of a discrete set of pose angles. This is performed in the context of face recognition using a database of face images of 21 people at nine discrete poses. In this formulation, pose estimation is performed first by computing the residual description error for a given image in the eigenspace for each pose. Recognition follows by finding the individual whose image is closest to the given image in the eigenspace for the selected pose. Results given indicate slightly better performance for this technique than for the universal eigenbasis, but these results are difficult to interpret due to the very coarse sampling in pose.

2.3.3 Classification of Commercial Aircraft

An application of orthogonal subspace projection methods to recognition of aircraft from range-profile data is presented by Zyweck [90, 91, 92]. In this work, range-profiles were collected from several aircraft as they took off from Adelaide Airport in Australia. Each aircraft was either a Boeing 727 or a Boeing 737, and range-profiles were collected over a 50 degree sweep in aspect angle.

The authors apply a number of preprocessing steps to the range-profiles prior to classification:

- computation of the magnitude of the return in each range bin,
- alignment of the range-profiles by correlation,

- averaging of 30 consecutive range-profiles to improve the signal-to-noise ratio,
- estimation of the noise level and subtraction of this value from each range bin, and
- computation of the magnitude of the Discrete Fourier Transform of the processed range-profiles.

The resulting preprocessed data are referred to as feature vectors. These are divided into training and testing sets. All of the feature vectors in the training set for a particular aircraft at all aspect angles are collected together and are used to compute an orthogonal basis using Fisher's Linear Discriminant. Because there are only two classes of aircraft, there is only a single basis vector, and the projection of a feature vector produces a scalar. Each of the feature vectors in the training set for a particular aircraft are projected into the orthogonal subspace, producing a set of scalars for each class. The mean and variance in each class are computed.

Recognition of a feature vector from the testing data set is performed by first projecting it into the orthogonal subspace. The resulting scalar is modeled as a Gaussian random variable, with mean and variance given by the class mean and variance computed from the training data. Classification of the projected feature is performed by a likelihood ratio test. The results of their test indicate a 96% rate of correct recognition was achieved for this two-class problem. This result is significant because a high rate of recognition was achieved using actual radar data and a very simple recognition algorithm. Furthermore, the algorithm achieved correct recognition for variations in aspect angle of 50 degrees. The authors indicate that a possible reason for their success lies in the different engine positions for the 727 and 737 aircraft. The engines are clearly visible to the radar over the range of aspect angles considered and are major scatterers, thus rendering the observed range-profiles significantly different for the two aircraft. It is difficult to say whether the techniques employed here would also be successful when attempting to distinguish a greater number of targets over a broader range of aspect angles.

2.4 Joint Tracking and Recognition

As discussed in Chapter 1, the problem of recognition of targets from sensor data is best solved by jointly estimating the position, orientation and target type from

the available data. In contrast, the vast majority of approaches to the recognition problem in the radar context assume that the orientation of the object is known, and ignore the problem of how orientation estimates might be obtained from the radar data. This section identifies those few efforts in which the authors recognize the advantages of jointly estimating the pertinent target parameters.

2.4.1 Joint Position and Orientation Estimation

Early work by Kendrick [39] and later advancements by Andrisani [3] and Swarder [82] seek to improve the accuracy of Kalman filter-based tracking algorithms by incorporating an estimate of one or more orientation angles derived from image data of the target. The relation between the position and orientation of an aircraft is nonlinear, but this relationship is apparently well behaved enough that the process of forward prediction and residual analysis is of significant utility in orientation-enhanced tracking. All three studies indicate superior performance of these tracking systems over traditional trackers which do not estimate orientation. Although none of these works considers identification directly, they represent important advancements of the view that the process of extracting information from a target using radar reflections is best performed by estimating all target parameters of interest, and that this can be achieved by processing data received from multiple sensors. A more detailed review of these works is provided by Libby [45].

2.4.2 Joint Tracking and Recognition by Sequence Comparison Methods

Libby [45, 44] provides a detailed comparison of six algorithms for orientation estimation and target identification within a general dynamic programming structure. A particular aircraft model is chosen, along with a maneuver which this hypothetical target will execute; the maneuver consists of true position and orientation states of the target over time. The sequence of orientation states is used to select a series of HRR signatures from a library of simulated returns for each potential aircraft on a discrete grid in orientation space. Next, range, range rate and pointing angle measurements are provided to a Kalman filter-based tracking algorithm. This algorithm provides estimates of the complete kinematic state of the target over time, including estimates of the orientation, which define an expected path for the target through

orientation space, and bounds on orientation errors, which are used to define windows around each estimate in which the true orientation is assumed to lie.

At this point, dynamic programming is applied to estimate the sequence of orientations corresponding to the sequence of observed HRR signatures. For each time step, the two most recent observations are matched to library elements drawn from the associated orientation windows to provide an estimate of a single transition in the orientation path. These transitions are connected to provide an estimated orientation sequence and the process is repeated for each potential target. The six algorithms considered are as follows:

- a Perfect Knowledge of Aspect algorithm in which library elements for each potential target are extracted for the true orientation sequence and compared with observations to compute the likelihood of each target having produced the observations;
- an Independent Look algorithm which finds the best match between each observation and the library elements from the associated orientation window without regard to the resulting path in orientation space;
- a Fixed Bound algorithm in which the only restriction on successive orientation estimates is that the change is bounded;
- an algorithm based on the work of Larson and Peschon in which the event of the target being in a particular orientation state at any time is assigned a probability conditioned on the previous orientation state, the entire history of kinematic measurements, and the target type;
- a pair of algorithms based on Dynamic Time Warping in one and two dimensions, respectively, in which transitions in orientation space are restricted to neighboring elements.

A number of experiments are reported where likelihoods of the true target and a different target having produced the observed sequence of measurements are computed. Libby seeks to analyze the performance of the classification algorithms in terms of generalized ambiguity functions. To accomplish this, he defines pseudo-targets lying between the two actual targets through the use of a morphing algorithm with various morphing fractions. The curvature of the resulting functions of likelihood

versus morphing fraction in the vicinity of the truth is qualitatively related to the notion of a Cramer-Rao Lower Bound on classification performance.

The results of these experiments show that, of those algorithms that were not provided with the true aspect angle sequence, the Larson and Peschon algorithm exhibited superior performance in terms of a lower computed likelihood for the incorrect target model than that computed by other algorithms. An interesting observation is that, for most of the cases considered, all six algorithms were able to correctly classify the target from the sequence of observations. Only in the cases of scatterer-augmented targets and high noise or the use of two target models which differed only in the presence of armaments under the wings did any of the six algorithms make an incorrect classification.

2.5 Other Methods Using Radar Data

The problem of recognition using radar has a long history in the literature, and a wide variety of methods have been proposed. To complete the review of ATR methods using range-profile data, this section outlines two efforts, one using correlation filters and another using wavelet representations, that would not be properly placed under the headings of previous sections. We conclude by presenting two examples in which the modulation introduced to the reflected signal by the propulsion system of an aircraft is used in identification.

2.5.1 Recognition Using Correlation Filters

Hudson [34] considers the use of a bank of correlation filters for recognition of aircraft. Normalized range-profiles for a number of aircraft are divided into groups, where each group contains the data for a particular aircraft over a 20 degree sweep in aspect angle. Each group of range-profiles is used to determine a vector of filter coefficients that maximizes the correlation peak between the filter vector and the range-profiles under all possible circular shifts of the range-profiles. This produces a correlation filter tuned to each target over a 20 degree window in aspect angle.

The resulting filter bank is tested using a database of real range-profiles collected from 24 different aircraft types in flight. Normalized range-profiles are divided into training and testing subsets. The training set is used to compute the filter coefficients, and the testing set are used to test the recognition performance of the

resulting filter bank. In the original test, an overall rate of correct recognition of 57% was achieved. This rate was deemed to be unacceptably low for practical use. The authors proceed to demonstrate that higher rates of recognition are achievable when the aspect angle is assumed known (79%), or when a recognition is performed using each of a sequence of eight observed range-profiles and the final estimate of target type for the sequence is determined by majority vote (84%). The authors do not comment on the use of the correlation filters to estimate the target aspect.

2.5.2 Recognition Using Wavelet Representations

Baras [5] develops efficient databases of high resolution radar data for use in ATR through the application of wavelet representations. Range-profiles for two ship targets simulated using a software package developed by the Naval Research Laboratory are used in the development and testing of the algorithm. The simulations are performed for aspect angles between 0 and 360 degrees in 0.5 degree increments. All of the range-profiles for both targets are processed using a wavelet multiresolution analysis to produce approximations to the data at four different scales, the finest scale corresponding to the original range-profiles.

The data for one of the ships at the coarsest scale are then grouped into cells in aspect angle over which the coarse-scale data are relatively invariant. A cell is defined as the range in aspect angle over which the mean squared distance between data vectors in the cell remains below a threshold. The coarse-scale cell having the largest mean-square distance is then subdivided into smaller cells, and the size of the sub-cells is determined using the mean-square distance between data vectors at the next finer scale. This process is repeated until the reduction in distortion by subdividing a given cell falls below a predetermined threshold. The result is a tree structure in which the leaf nodes represent regions in aspect angle over which the data vectors at a particular scale are consistent with respect to the mean-square distance metric. This tree is known as an aspect graph.

Recognition is performed using the resulting aspect graphs. A received range-profile is preprocessed using wavelet analysis as described earlier. The resulting data vectors are input to the aspect graph. At each node, the distances between the sensor data and the data for each node at the next lower level are computed, and the sub-node corresponding to the smallest distance is selected. This is repeated until a leaf node is reached, at which time a classification is assigned according to the label

assigned to the leaf node. The authors investigate two methods for implementing this structure, one in which a single aspect graph is developed from the data for both ships, and another in which separate aspect graphs are developed for each ship. Not surprisingly, the method using separate aspect graphs for each ship achieved a higher recognition rate (85%) than did the method using a single aspect graph (75%). This effort is important because it provides a technique for nonuniform sampling of the possible aspect angles, where the sampling is driven by the variability in the data. It is interesting to note that while the algorithm simultaneously provides estimates of the target type and aspect angle, no results are reported with respect to the accuracy of aspect angle estimates obtained.

2.5.3 Modulation by Propulsion Substructure

Bell [9] describes a system in which distortion of the radar reflection by an aircraft's propulsion system is treated as target-specific information to be used in identification. The periodic rotation of jet engines is modeled as introducing periodic amplitude and phase modulation to the received waveform at integer multiples of the rotation frequency. This distortion of the received signal is known as jet engine modulation or JEM. Fourier representations of the modulation functions are estimated by locating peaks in the periodogram spectra and cepstra of radar reflections. An example is shown where the number of blades on a dual propeller aircraft is derived from the harmonic relationship between the periodogram peaks, but no results are shown where this technique is used for identification.

Bullard [15] proposes Doppler processing of radar returns from a helicopter target to estimate parameters of the rotor structure which are indicative of the type of helicopter in question. Responses from rotors are assumed to lie in regions in the frequency domain distinct from responses from other parts of the aircraft, allowing selection of the rotor responses by simple filtering. The time delay between peaks in the rotor response is used to estimate the rotation rate. It is claimed that many other rotor parameters can be extracted using similar processing, but details are not given. No identification studies are reported.

Work on JEM is ongoing. The general result is that JEM information is useful for recognition, but requires long dwell times to acquire, and may be impossible to measure if the orientation is such that the target's propulsion system is not directly illuminated.

2.6 Radar Reflectivity Models

As discussed in Chapter 1, recognition of targets from their radar reflections requires mathematical models that accurately characterize the data expected from a given target at a given orientation. The process by which radar signals reflect off a complex surface to produce an echo signal can be modeled in a variety of ways. The most general model for this phenomenon would consider the complete interaction of electromagnetic radiation with the target, including reflected and absorbed waves, induced surface currents and propagation effects. The complicated nature of such a model would require a prohibitive amount of computation to understand the behavior of radar reflections for all but the simplest of targets. As a result, most practical models of radar reflectivity are simplified to accommodate the needs of the modeler (e.g. simulation). This section provides a brief list of several representative models.

2.6.1 One-Dimensional Models

This section describes two approaches to modeling the range-profile as the summed contribution of a small number of scattering centers. I have referred to these models as one-dimensional because the dependence of the return from any scattering center on position is only in terms of the range coordinate. These models are generally used to achieve a dimensional reduction in existing range-profile data, and cannot be used to compute the range-profile from, for example, a 3-D description of the target geometry.

Impulse Train

The simplest model for HRR data is provided by Smith [73]. Drawing on the widely held view that the structure of a range-profile for a particular target is dominated by the contributions of a small number of scattering centers, Smith models each of these using a scaled impulse function,

$$r(t) = \sum_{l=1}^m A_l \delta(t - t_l). \quad (2.3)$$

Thus, each scatterer is parameterized only by its range and amplitude. Smith points out that a scattering center does not necessarily correspond to an actual physical structure on the target, and that the modeling of scatterers as impulses may preclude

the use of this technique in describing returns from cavities and ducts. A single example is reported in which a range-profile is computed from a facet model for the Lockheed SR-71 using an unspecified simulator, and the parameters corresponding to 26 scattering centers are estimated. These parameters are used to reconstruct an estimated range-profile which appears very similar to the original data, although no quantitative analysis is provided.

Prony Model

A widely used scattering center model for range-profiles which includes polarization effects is known as Prony's model and is described by Steedly [79]. In this case, the horizontally and vertically polarized responses to a left-circularly polarized transmission are characterized as being dominated by the contributions of M scattering centers, each of which constitutes a pole of the model:

$$\begin{bmatrix} s_{hl}(f_n) \\ s_{vl}(f_n) \end{bmatrix} = \sum_{k=1}^M \begin{bmatrix} a_{hk} \\ a_{vk} \end{bmatrix} p_k^n, \quad n = 1, \dots, N. \quad (2.4)$$

The phase of each p_k is proportional to the range of the associated scattering center, and the complex coefficients a_{hk} and a_{vk} describe the polarization properties. These properties are commonly described in terms of a polarization ellipse for each scatterer, which is in turn parametrized by its amplitude, ellipticity and tilt angle. Steedly proposes estimation of these parameters using linear prediction and least squares, and applies this technique to data collected from scaled aircraft models in a compact radar range. Unlike Smith, Steedly makes a strong connection between the estimated scattering centers and physical structures on the target; accurate estimation of the range and polarization parameters for each scatterer is used to validate the model and estimation procedure.

2.6.2 Three-Dimensional Models

Two efforts are presented in which the range-profile is modeled as a deterministic quantity, and algorithms have been developed for computation of the range-profile from a 3-D characterization of the target shape. These models can be used to understand the variation in the range-profile with orientation. Simulated range-profiles are computed for a discrete set of orientations, and appropriate distance metrics or statistical measures are computed between pairs of range profiles.

Single Reflection with Shading

Du [25] describes a system for range-profile prediction for a target whose surface is constructed of quadrilateral patches. As this effort was intended to be of use to the HRR group at Rome Laboratory, the system is designed to simulate the output of existing radar hardware, which is the convolution of the actual range-profile with the product of a transmitted chirp waveform and a sinc function.

The process of range-profile prediction proceeds in two stages. First, for a particular target and orientation, a shading mask is computed which consists of those patches that are directly illuminated by the radar system. Illuminated patches are subdivided into four corners and the return in each range bin is the summed contribution of all patch corners in the bin. The return from each corner is the product of a random phase factor and reflection factor which is an appropriately chosen function of the corner area and the observation angle. The particular function chosen for use in the simulator is the antenna pattern of an appropriately sized and oriented circular aperture:

$$f(\xi, A) = A \frac{J_1(2\pi \sin(\xi)r/\lambda)}{2\pi \sin(\xi)r/\lambda}, \quad (2.5)$$

where ξ is the angle between the radar line-of-sight and the corner normal vector, A is the area both of the corner and the associated circular aperture, and r is the aperture radius. Du goes on to describe several other possible choices for this function. Computer code which implements the algorithm is provided and several issues relating to implementation are discussed.

XPATCH

XPATCH is a set of four computer programs used to simulate the radar return of a target [2]. A set of parallel rays are directed toward a faceted model of the target, and traced as they bounce off the various surfaces. Those rays which are eventually directed back toward the source contribute to the return. This shooting and bouncing ray technique accounts both for shading and multiple bounce effects. The output of the XPATCH programs is either radar cross section as a function of aspect angle, the range-profile for a particular aspect angle, or an ISAR image. For the range-profile case, the return is computed for each of the four combinations of linear transmit and receive polarizations. A more complete description of XPATCH is provided in Chapter 3.

2.6.3 Stochastic Models

The following efforts model the interaction of electromagnetic radiation with the target as the realization of a random process. One model has been analyzed to yield the expected variation of the complete return from the target with orientation.

The WSSUS Model

Van Trees [87] presents a statistical model for the reflectivity of a doubly-spread radar target whose surface is rough on the scale of the wavelength of the transmitted signal, and is rotating about a point on the radar line-of-sight. The reflectivity of the target is denoted by $b(t, \tau)$, and is modeled as a zero-mean complex Gaussian random process. Furthermore, it is assumed that the reflectivities at different delays τ are statistically uncorrelated, and that for a fixed delay τ the reflectivity is wide-sense stationary in time. A derivation of the WSSUS model is presented in Chapter 3.

Conditionally Gaussian Model

In this model, the received data are modeled as the sum of a signal term and white Gaussian noise. The signal term is modeled as a random process that is conditionally Gaussian given the target type and orientation. A derivation of this model and a number of results with respect to its use in recognition are presented in Chapter 3. As described earlier, the Adaptive Gaussian Classifier models the processed radar data as conditionally Gaussian, in the sense that classification is performed using the Gaussian log-likelihood function.

Laser Radar Model

Shapiro [68, 69, 70] considers illumination of a target by the focused beam of a coherent laser radar. The beam radius is assumed to be small on the scale of the target and an image is formed by making separate measurements of the target reflectivity for different transmitter aim points. High resolution range determination for each aim point allows for a partial 3-D characterization of the target.

Under the conditions of a stationary unresolved target and a monostatic radar system, the complex envelopes of the incident and reflected fields $\bar{\mathbf{E}}_i$ and $\bar{\mathbf{E}}_r$ are related through a multiplicative target model

$$\bar{\mathbf{E}}_r(\bar{\rho}, t) = \mathbf{T}(\bar{\rho})\bar{\mathbf{E}}_i(\bar{\rho}, t), \quad (2.6)$$

where $\bar{\rho}$ is the transmitter aim point in the image plane and $\mathbf{T}(\bar{\rho})$ is the associated complex-field reflection coefficient. This quantity is decomposed into a deterministic component $\mathbf{T}_g(\bar{\rho})$ which models the target as a specular reflector such as a polished surface and a random component $\mathbf{T}_s(\bar{\rho})$ which characterizes the effect of the roughness of the target surface

$$\mathbf{T}(\bar{\rho}) = \mathbf{T}_g(\bar{\rho})e^{j\theta} + \mathbf{T}_s(\bar{\rho}). \quad (2.7)$$

The glint component $\mathbf{T}_g(\bar{\rho})$ is considered to have a very narrow angular extent like that of a mirror, but any appropriate reflection pattern could be substituted for this term. The phase angle θ is modeled as a uniform random variable on $[0, 2\pi]$. The speckle component $\mathbf{T}_s(\bar{\rho})$ is a complex-valued 0-mean Gaussian random process with autocorrelation function

$$E \{ \mathbf{T}_s(\bar{\rho}_1) \mathbf{T}_s(\bar{\rho}_2) \} = 0, \quad (2.8)$$

$$E \{ \mathbf{T}_s(\bar{\rho}_1) \mathbf{T}_s^*(\bar{\rho}_2) \} = \lambda^2 \mathcal{T}_s(\bar{\rho}_1) \delta(\bar{\rho}_1 - \bar{\rho}_2). \quad (2.9)$$

Shapiro proceeds to provide a characterization of the directional properties of the speckle component [70]. In this case, $\mathbf{T}_g(\bar{\rho})$ is ignored and the multiplicative model is expressly parameterized in terms of the position and orientation of the target

$$\mathbf{T}(\bar{\rho}) = \mathbf{T}_s \left[\bar{\mathbf{U}}_\theta(\bar{\rho} - \bar{\rho}_c) \right] \exp \left[j2k\bar{\psi}^T(\bar{\rho} - \bar{\rho}_c) \right], \quad (2.10)$$

where $\bar{\rho}_c$ is a target translation vector, $\bar{\mathbf{U}}_\theta$ is a matrix describing target rotation in the plane of illumination and $\bar{\psi}$ is a vector describing the tilt of the target with respect to the plane of illumination. Note that object tilt is equivalent to a change in aspect angle. The complex amplitude of the return due to illumination by a monochromatic beam of spatial pattern $\xi_L(\bar{\rho})$ is given by

$$\mathbf{y} = \int \mathbf{T}(\bar{\rho}) \xi_L^2(\bar{\rho}) d\bar{\rho}. \quad (2.11)$$

The variation in target reflectivity is investigated by computing the correlation coefficient γ_{12} of the random variables $|\mathbf{y}_1|^2$ and $|\mathbf{y}_2|^2$, where the subscripts denote different choices of the parameters $\{\bar{\rho}_c, \bar{\mathbf{U}}_\theta, \bar{\psi}\}$. If $\mathbf{T}(\bar{\rho})$ is modeled as a complex Gaussian random process, then we have

$$\gamma_{12} = \frac{|E \{ \mathbf{y}_1 \mathbf{y}_2^* \}|^2}{E \{ |\mathbf{y}_1|^2 \} E \{ |\mathbf{y}_2|^2 \}}. \quad (2.12)$$

Speckle decorrelation is said to occur when the measurement parameters for the two different observations are such that $\gamma_{12} \leq e^{-2}$. For the specific case of object tilt $\bar{\psi}_1 - \bar{\psi}_2 = \Delta\bar{\psi}$, decorrelation occurs for

$$|\Delta\bar{\psi}| \geq \frac{2a}{\sqrt{4L^2 + k^2a^4}} \quad (2.13)$$

where a is the transmitted beam radius, L is the distance to the target and k is the wave number of the illumination. When only tilt is varied, the decorrelation limit is approximately $2/ka$ for near-field illumination and a/L for far-field illumination. If multiple parameters are allowed to vary, the decorrelation limit is $2/ka$, regardless of the value of L .

Shapiro's model has yielded a measure of the expected variation in the reflectivity of a target for small changes in orientation, when the signal-target interaction can be modeled as taking place in a plane. In order to more directly relate this result to the case of an extended target on which all points are simultaneously illuminated by a wide radar beam, we will model the beam pattern $\xi_L(\bar{\rho})$ with a constant value of ξ over the entire image plane. If we additionally ignore translation and in-plane rotation, then we have

$$E\{y_1 y_2^*\} = \lambda^2 \xi^4 \int d\bar{\rho} T_s(\bar{\rho}) e^{j2k(\bar{\psi}_1 - \bar{\psi}_2)^T \bar{\rho}}. \quad (2.14)$$

2.7 Conclusions

In this chapter, we have reviewed a wide variety of approaches to the problem of recognition from radar data, and identified several models for high resolution radar data that have been applied to the recognition problem. Of the models examined, only the laser radar model described by Shapiro has been investigated to determine the variation in the data for small changes in orientation. None have been analyzed in terms of fundamental limits on the performance of algorithms for orientation estimation and recognition. We will further examine the modelling problem and analyze two candidate models in Chapters 3 and 4.

Those algorithm for ATR that have been tested on real or simulated data from realistic radar targets have met with varying degrees of success, but because of the vast differences in the methods by which these algorithms were tested, it is difficult to draw comparisons between them. However, it is important to note the distinctions

between these efforts and the approach employed in this thesis. Very few of the efforts pose the recognition problem in a Bayesian framework. Of the various ATR algorithms, only the Adaptive Gaussian Classifier makes use of a stochastic model for the range-profile data. Only the effort by Libby attempts to jointly estimate the target type and a sequence of target orientations from a sequence of observed range-profiles, but this method does not relate the successive orientations assumed by the target through a dynamical model for the target motion. In short, no previous effort we have found incorporates all of the essential elements of the recognition algorithms we describe in Chapter 5.

Chapter 3

High Resolution Radar Models

Development of systems for recognition of aircraft and ground vehicles from high resolution radar data is critically dependent on the use of appropriate models that characterize the range-profiles expected from a given target at various orientations. In this chapter we consider the HRR modeling problem. We begin by introducing the reader to various basic concepts in high resolution radar, drawn largely from standard textbooks [72, 88, 64, 87], including a definition of the basic unit of high resolution radar data that will be used for recognition of targets, the range-profile. We go on to consider two candidate models for the range-profile, and conduct extensive tests of the recognition performance under these models using simulated radar data. We conclude by indicating additional work to be performed to further establish the utility of models for high resolution radar data in performing recognition.

3.1 Basic Concepts in Radar Modeling

A fundamental concept associated with radar systems is that range is proportional to delay. For example, the range to an object being illuminated by the radar is proportional to the time delay between transmission of a radar pulse and reception of the echo pulse. The range R is related to the delay τ through [88]

$$R = \frac{\tau c}{2}, \tag{3.1}$$

where c is the velocity of propagation of the signal wavefront. As will be shown, the equivalence of range and delay can be exploited both to determine the bulk range

to a target, and to infer the distribution of target reflectivity as a function of range along the radar line-of-sight.

A second important concept in radar is that the ratio of power in the received echo to power in the transmitted pulse is proportional to a quantity called the radar cross-section or RCS. This relationship is summarized in the radar range equation, which states that for a target at range R , the transmitted power P_T and received power P_R are related by [88]

$$P_R = P_T G_T \left(\frac{1}{4\pi R^2} \right) \sigma \left(\frac{1}{4\pi R^2} \right) A_R, \quad (3.2)$$

where σ is the radar cross-section of the target, G_T is the gain in the transmitter antenna, A_R is the effective aperture area of the receive antenna, and the two factors of $(1/4\pi R^2)$ describe the effect of the spreading of the signal wavefront as it travels from the radar to the target and back. Thus, if the target is motionless, and if the radar parameters P_T , G_T , and A_R are fixed, we have that

$$P_R \propto \sigma. \quad (3.3)$$

This provides a simple mechanism to measure the radar cross-section of the target. Examination of equation (3.2) reveals that the radar cross-section is expressed in units of square meters; this quantity can be thought of as the effective echoing area of the target.

3.1.1 Point Targets

To begin the development of models for radar signals, consider the target to be a zero-velocity isotropic point reflector. Such a target is considered to have negligible size, such that the reflection of the impinging wavefront occurs instantaneously. Let $\sqrt{E_T} s_T(t)$ and $r(t)$ be the complex envelopes of the transmitted and received signals, respectively. In this characterization, the transmitted signal has energy E_T , and $s_T(t)$ is assumed normalized to unit energy. Let T be the duration of the transmitted signal, so that the transmitted power is given by

$$P_T = \frac{E_T}{T}, \quad (3.4)$$

and is expressed in units of Watts. Let the target be located at two-way delay τ from the radar transmitter and have reflectivity b . Then the received signal is modeled as a scaled, delayed version of the transmitted signal, with additive noise [87],

$$r(t) = g(R)b\sqrt{E_T}s_T(t - \tau) + w(t). \quad (3.5)$$

In this model, the reflectivity b includes only the effects of reflection by the target; if b is deterministic, then its squared magnitude gives the radar cross-section in units of square meters,

$$|b|^2 = \sigma. \quad (3.6)$$

The factor $g(R)$ includes all of the range attenuation and antenna effects,

$$g(R) = \frac{\sqrt{G_T A_R}}{4\pi R^2} \quad (3.7)$$

and has units of meters^{-1} . The noise $w(t)$ is a zero-mean complex white Gaussian random process of intensity N_0 Watts per Hertz.

In the absence of noise, $N_0 = 0$, the received power in Watts is given by

$$P_T = \frac{1}{T} \int_{\tau}^{\tau+T} |r(t)|^2 dt. \quad (3.8)$$

For consistency with the radar range equation (3.2), we have

$$P_T = g^2(R)|b|^2 \frac{E_T}{T}. \quad (3.9)$$

For situations where the target is small, or the radar bandwidth is low, this model can be used to determine the bulk range to the target, or if the range is known, to measure the radar cross-section σ [72, 56].

It is typically the case that an object being illuminated by a radar is in motion. If the point target is allowed to move with radial velocity v along the line-of-sight, then the range R and delay τ become functions of time, and this must be reflected in the model. Let the transmitted pulse have bandwidth W . As shown in [87], if the target velocity is small compared to c ,

$$\frac{v}{c} \ll 1, \quad (3.10)$$

and the time-bandwidth product for the transmitted signal satisfies

$$WT \ll \frac{c}{2v}, \quad (3.11)$$

then the effect of the motion on the received signal is to introduce a Doppler shift to the carrier frequency ν_c , other effects being negligible. Thus, the complex envelope of the received signal becomes

$$r(t) = g(R)b\sqrt{E_T}s_T(t - \tau)e^{j2\pi\nu_D t} + w(t), \quad (3.12)$$

where

$$\nu_D = \nu_c \left(\frac{c}{2v} \right) \quad (3.13)$$

is the Doppler frequency.

3.1.2 Range Resolution and Range-Spread Targets

The range resolution of a radar system is defined by the ability to distinguish point reflectors closely spaced in range. For a transmitted signal of bandwidth W , two reflectors are distinguishable if they are separated in range by no less than [88]

$$\Delta R = \frac{c}{2W}. \quad (3.14)$$

High range resolution can be defined as the case where the target extent in the range dimension, L , is much larger than the range resolution

$$L \gg \Delta R. \quad (3.15)$$

In this case the target is said to be range-spread or dispersive. If we continue to model the target as a reflector, and if the target is motionless, then the received signal is given by the superposition of reflections from all delays τ over the target extent [87],

$$r(t) = g(R)\sqrt{E_T} \int_0^{\frac{2L}{c}} b(\tau)s_T(t - \tau) d\tau + w(t) \quad (3.16)$$

where $b(\tau)$ is, in turn, the superposition of reflectivities for all points on the target at delay τ . In (3.16), it is assumed that the bulk range to the target R is much larger




Figure 3.1: Illustration of high resolution radar. The target is illuminated by a short pulse $s_T(t)$. This signal is reflected by the target with reflectivity profile $b(\tau)$ to produce the received signal $r(t)$ according to (3.16). The reflectivity $b(\tau)$ is the superposition of reflectivities of all points on the target at two-way delay τ , and is dependent on the orientation of the target relative to the sensor.

than the target extent L , such that the factor $g(R)$ is approximately constant with respect to the integration.

The scenario for collection of high resolution radar data from a range-spread target is illustrated in Figure 3.1. The target is illuminated by a short pulse $s_T(t)$. This signal is reflected by the target with reflectivity profile $b(\tau)$ to produce the received signal $r(t)$ according to (3.16). Thus, for a range-spread target, the received signal $r(t)$ contains information about the distribution of target reflectivity with range. It is assumed that the beamwidth of the transmitted signal is significantly larger than the target size, and that the radar antenna is aimed at the target so that the target lies within the main lobe of the antenna pattern. This scenario for radar illumination is known as spotlight mode. Under these assumptions, the received signal $r(t)$ is invariant to small translations of the target in any direction normal to the radar line-of-sight, i.e. translations in the cross-range dimension. However, $r(t)$ is clearly dependent on the target range, and also on physical factors such as

- the materials from which the target is constructed,
- the shape of the target, and
- the orientation of the target relative to the radar sensor.

It is the last two of these items that we will exploit in using HRR data to identify targets.

When collecting high resolution radar data according to the scenario described above, it is important to note that there is typically some uncertainty in the measurement of the bulk range to the target. As a result, it is customary to include a phase factor $e^{j\phi}$, where ϕ is a random variable, multiplying the signal term in the reflectivity model (3.16). If the uncertainty in the bulk range is larger than a wavelength of the transmitted signal, then ϕ is well modeled as uniform on the circle. There are various ways to account for this random phase factor in our model, one of which would be to include it in $b(\tau)$.

3.1.3 Doubly-Spread Targets

Here we consider the case where a range-spread target is in motion. If the motion is not strictly directed along the line-of-sight, then the orientation of the target relative to the radar will change as the target moves. Thus the reflectivity profile b will change with time. Let $b(t, \tau)$ represent the superposition of reflectivities for all points on the target at delay τ and for an illumination occurring at time t . Let the radar transmit a pulse toward the target at time $t - \tau$. The signal arrives at a point on the target with two-way delay τ at time $t - \tau/2$. The combined reflectivity of all such points is [87]

$$b\left(t - \frac{\tau}{2}, \tau\right).$$

The reflected signal arrives at the radar at time t . In this case, the received signal is given by [87]

$$r(t) = g(R)\sqrt{E_T} \int_0^{\frac{2L}{c}} s_T(t - \tau) b\left(t - \frac{\tau}{2}, \tau\right) d\tau + w(t) \quad (3.17)$$

For general target motion, there will be a component of the velocity that is rotational about the target centroid. Points on the target at varying distances y from the line-of-sight will have a radial component of their velocity that is proportional to y . Thus, the Doppler shift introduced to the reflected signal by a point at cross-range

coordinate y will be proportional to y . It is this dependence that is exploited in radar imaging applications such as SAR and ISAR. If $b(t, \tau)$ is modeled as a deterministic function of time and delay, then let $c(\nu, \tau)$ be the Fourier transform of $b(t, \tau)$ [74],

$$c(\nu, \tau) = \int_{-\infty}^{\infty} b(t, \tau) e^{j2\pi\nu t} dt. \quad (3.18)$$

The function $c(\nu, \tau)$ is the superposition of reflectivities for all points on the target with range coordinate corresponding to delay τ and with cross-range coordinate corresponding to Doppler shift ν .

For distributed targets with many scatterers, it is often appropriate to model the reflectivity b as a random process in time and delay. As described in [87], if the target surface is rough compared to the wavelength at the carrier frequency, then $b(t, \tau)$ may be modeled as a zero-mean complex Gaussian random process with covariance function

$$E\{b(t_1, \tau_1)b^*(t_2, \tau_2)\} = K(t_1 - t_2, \tau_1)\delta(\tau_1 - \tau_2). \quad (3.19)$$

The form of the right side of this equation arises from two assumptions made about the reflectivity process:

- The reflectivities at different delays τ are statistically independent, and
- The reflectivity at a fixed delay forms a wide-sense stationary random process.

These assumptions form the basis for the wide-sense stationary, uncorrelated scatterers or WSSUS model for doubly-spread radar targets [87]. Assuming this model, let $S(\nu, \tau)$ be the Fourier transform of the covariance function $K(t, \tau)$,

$$S(\nu, \tau) = \int_{-\infty}^{\infty} K(t, \tau) e^{-j2\pi\nu t} dt. \quad (3.20)$$

This function is known as the scattering function for the target.

It is straightforward to show that, in the absence of noise, the radar cross-section of a doubly-spread target is proportional to the integral of the scattering function over the target extent in delay and Doppler. Let

$$r(t) = s_R(t) + w(t) \quad (3.21)$$

so that $s_R(t)$ is the signal portion of the received waveform that results from reflection. The expected received power is given by

$$\begin{aligned}
E\{P_R\} &= \frac{1}{T} \int_T E\{|s_R(t)|^2\} dt \\
&= \frac{E_T}{T} g^2(R) \int_T dt \int \int E\{b(t - \frac{\tau}{2}, \tau) b(t - \frac{\sigma}{2}, \sigma)\} s_T(t - \tau) s_T(t - \sigma) d\tau d\sigma \\
&= P_T g^2(R) \int K(0, \tau) \int_T |s_T(t - \tau)|^2 dt d\tau \\
&= P_T g^2(R) \int \int S(\nu, \tau) d\nu d\tau
\end{aligned}$$

Therefore, we have

$$\sigma = \frac{E\{P_R\}}{P_T g^2(R)} = \int \int S(\nu, \tau) d\nu d\tau. \quad (3.22)$$

Given this relationship, the RCS of a small patch on the target consisting of all points with delay coordinate in $[\tau, \tau + \Delta\tau)$ and Doppler coordinate in $[\nu, \nu + \Delta\nu)$ is proportional to $S(\nu, \tau) \Delta\nu \Delta\tau$ [74].

3.1.4 The Range Profile

The basic unit of high resolution radar data that will be used for recognition of targets is the range-profile. In this subsection, a derivation of the range-profile is presented for the range-spread target model. The development presented is drawn largely from the text by Rihaczek [64], with some notational changes for consistency with the WSSUS model.

The derivation is begun by examining the signal term in equation (3.17). We assume that the duration of the coherent target interaction as prescribed by this equation is short enough that target motion during a single illumination is negligible. This is equivalent to assuming that the target is unresolved in the Doppler coordinate. In this case, equation (3.19) may be written

$$E\{b(\tau_1) b(\tau_2)\} = K(\tau_1) \delta(\tau_1 - \tau_2), \quad (3.23)$$

where $K(\tau) = E\{|b(\tau)|^2\}$ is the range-scattering function for the target. Let ν_0 be the bulk Doppler shift resulting from the target radial velocity. In the absence of noise, the received signal is

$$s_R(t) = g(R) \sqrt{E_T} \int_{-\infty}^{\infty} s_T(t - \tau) e^{j2\pi\nu_0 t} b(\tau) d\tau. \quad (3.24)$$

Let τ_0 be the bulk delay to the target. The received signal is passed through a filter matched to the transmitted signal delayed by τ_0 and Doppler shifted by ν_0 [64],

$$\psi_F(t) = s_T(t - \tau_0)e^{j2\pi\nu_0(t-\tau_0)}. \quad (3.25)$$

The matched filter impulse response is given by

$$h(t) = \frac{1}{g(R)\sqrt{E_T}}\psi_F^*(-t) = \frac{1}{g(R)\sqrt{E_T}}s_T^*(-t - \tau_0)e^{j2\pi\nu_0(t+\tau_0)}. \quad (3.26)$$

Let $\eta(t)$ be the output of the matched filter. Then we have

$$\eta(t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} s_T(\sigma - \tau)e^{j2\pi\nu_0\sigma} s_T^*(-t + \sigma - \tau_0)b(\tau)e^{j2\pi\nu_0(t-\sigma+\tau_0)} d\tau d\sigma \quad (3.27)$$

Making the substitution $\lambda = \sigma - \tau$, the expression for the matched filter output becomes

$$\eta(t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} s_T(\lambda)s_T^*(\lambda - t + \tau - \tau_0)b(\tau)e^{j2\pi\nu_0(t+\tau_0)} d\tau d\lambda. \quad (3.28)$$

Let the ambiguity function $\chi(\tau, \nu)$ for the transmitted signal $s_T(t)$ be given by [64]

$$\chi(\tau, \nu) = \int_{-\infty}^{\infty} s_T(t)s_T^*(t - \tau)e^{j2\pi\nu t} dt. \quad (3.29)$$

Then we have

$$\eta(t) = \int_{-\infty}^{\infty} b(\tau)\chi(t - \tau + \tau_0, 0)e^{j2\pi\nu_0(t+\tau_0)} d\tau \quad (3.30)$$

Finally, computing the squared magnitude of $\eta(t)$ and taking the expected value, we have

$$E\{|\eta(t)|^2\} = \int_{-\infty}^{\infty} K(\tau)|\chi(t - \tau + \tau_0, 0)|^2 d\tau. \quad (3.31)$$

Note that the ambiguity function χ evaluated at $\nu = 0$ is simply the autocorrelation function for $s_T(t)$. Thus, this derivation shows that the expected value of the squared magnitude of the matched filter output is equal to the range-scattering function $K(\tau)$ smeared out by the squared magnitude of the autocorrelation function for the transmitted signal. The signal $\eta(t)$, after possible further processing depending upon the details of the radar system in question, is referred to as a range-profile. Because of the strong relationship between the signal $\eta(t)$ as a function of time and the reflectivity $b(\tau)$ as a function of range, it is customary to ignore the spreading

introduced by the autocorrelation function for the transmitted signal, and to consider the received signal as a measurement of target reflectivity. Thus the range-profile is defined as a time signal, produced by applying matched filtering and possibly other processing to a received radar echo, and which is considered to be a measurement of the distribution of the target reflectivity along the radar line-of-sight. By the discussion in Section 3.1.2, the range profile is invariant to translation in cross-range, but is dependent on the range to the target, the orientation of the target relative to the sensor, and the shape of the target. The next section is devoted to describing appropriate models for the range-profile, with specific attention to the dependence on the distribution of reflectivity over the target surface and on the target orientation.

3.2 Modeling Range-Profiles

The previous section introduced a statistical model for radar data known as the WSSUS model. For targets unresolved in Doppler, this model describes the reflectivity profile $b(\tau)$ as a 0-mean Gaussian random process that is white in range or delay. The model was used to derive an expression for the basic unit of HRR data, the range-profile. While not made explicit in the notation, it is clear that $b(\tau)$ and therefore the range-profile will depend on the target orientation θ and on the target type a . The potential for joint estimation of $\{\theta, a\}$ from range-profile data under a given model is determined by the likelihood function for observing a particular range-profile, given the orientation and target type. In this section, we consider both deterministic and stochastic models for the range-profile, introduce candidate models, examine the likelihood functions under these models, and discuss their relative merits.

A general form of a model for the range-profile is given by

$$r(t) = s(t; \theta, a) + w(t), \quad (3.32)$$

where $r(t)$ is the complex envelope of the observed range-profile, $s(t; \theta, a)$ is the signal component of the waveform, and $w(t)$ is a white Gaussian noise process of intensity N_0 . Note that we have made a significant change of notation here from the previous section. The signal $r(t)$ no longer represents the received radar echo, but rather represents the output of the matched filter that in the previous section was referred to as a range-profile and denoted by $\eta(t)$. An actual radar system, depending on the exact form of the transmitted signal and other system-specific parameters, may apply

significant processing in addition to the matched filter to produce range-profiles. In the remainder of this section, we will think of $r(t)$ as the final processed output of the receiver, refer to it as a range-profile, and assume that it obeys the model (3.32). The signal component s depends on the orientation of the target, θ , and on the target type a .

We are of course interested in applying the models we develop to real or simulated radar data, which is generally not available as a continuous waveform, but as a sequence of samples of such a waveform. Therefore, we will consider a discretized version of the model (3.32). Let the sampling be performed by integration of $r(t)$ against a unit-energy rectangular pulse of duration equal to the sampling interval. The sampling interval is typically chosen to correspond to the range resolution of the radar system in question. The pulse is shifted to produce the sample in each range bin, and the resulting samples are collected into a vector \mathbf{r} . The discretized version of (3.32) becomes

$$\mathbf{r} = \mathbf{s}(\theta, a) + \mathbf{w}, \quad (3.33)$$

where \mathbf{s} is the signal vector, and \mathbf{w} is the noise vector. The entries of \mathbf{w} are a sequence of independent, Gaussian random variables with mean 0 and variance N_0 whenever the samples are collected by integrating against orthonormal functions. When possible, range-profile data are typically calibrated so that the squared magnitude of the return in each range bin represents a measurement of the radar cross section of that portion of the target. For consistency with this practice, we will express the squared magnitude of any element of \mathbf{r} and the noise variance N_0 in units of square meters.

Our interest in the HRR models we examine will be to evaluate the potential for recognition of targets from range-profile data. A primary factor making recognition difficult is the well-known fact that range-profiles exhibit extreme variability for small changes in target orientation, an effect known as speckle or scintillation. The mechanism responsible for speckle in radar returns, as reported variously in [72, 70, 21], is typically described in terms of phase cancellation. Consider monochromatic illumination of a radar target, or perhaps only that portion which lies within a particular range bin, by a source of wavelength λ . The reflected field is the sum of reflections from all points on the target. If the orientation of the target is changed so that points on the target move through a distance on the order of λ , then the relative phases contributing to the complete return will produce varying degrees of constructive and destructive interference over the rotation. If the target is large compared to λ , then the rotation required to produce this effect will be very small. An important factor

in evaluating HRR models will be the potential for recognition that is robust with respect to radar speckle.

Models of the form (3.32) may be divided into two broad categories depending upon whether the mapping from $\{\theta, a\}$ to $s(t; \theta, a)$ is deterministic or stochastic. The question of which model is appropriate is an open one, and will be addressed later in this section and in the next chapter. We begin by introducing a deterministic and a stochastic model for consideration.

3.2.1 Deterministic Model

The idea that the reflection of a radar signal by a macroscopic target is characterized by a deterministic process is intuitively appealing. Given a complete characterization of the target shape and the materials from which it is constructed, the methods employed by the electromagnetics community can be used to describe all of the interactions that take place between the target and the impinging wavefront. Thus it has been the viewpoint of many investigators that the range-profile is computable given the target orientation. In this case, given a and θ , $s(t; \theta, a)$ is deterministic, and the likelihood function for the range-profile may be derived in a straightforward manner using the results of [86].

Let the received signal $r(t)$ be written as the sum

$$r(t) = \lim_{M \rightarrow \infty} \sum_{m=1}^M r_m \phi_m(t), \quad (3.34)$$

where $\{\phi_m(t)\}$ is a set of basis functions that form a complete orthonormal set, the coefficients r_m are given by

$$r_m = \int r(t) \phi_m(t) dt \quad (3.35)$$

$$= \int s(t; \theta) \phi_m(t) dt + \int w(t) \phi_m(t) dt \quad (3.36)$$

$$= s_m(\theta) + w_m, \quad (3.37)$$

and convergence of the limit is in the mean square sense. The truncation of the series representation for $r(t)$ to M terms is denoted $r_M(t)$. The coefficients r_m are a set of independent Gaussian random variables with mean $s_m(\theta)$ and variance N_0 . Thus,

the likelihood function for $r_M(t)$ is given by

$$p(r_M|\theta) = \prod_{m=1}^M \frac{1}{\pi N_0} \exp \left(-\frac{1}{N_0} \text{Re} \{ [r_m - s_m(\theta)][r_m - s_m(\theta)]^* \} \right). \quad (3.38)$$

The likelihood function for the received waveform $r(t)$ is defined as

$$p(r|\theta) = \lim_{M \rightarrow \infty} \frac{p(r_M|\theta)}{p(r_M|s(t;\theta) = 0)}. \quad (3.39)$$

Substituting (3.38) into the numerator and denominator of (3.39), evaluating the limit as $M \rightarrow \infty$, and taking the logarithm yields the log-likelihood function for observing $r(t)$ conditioned on the target type a and the target orientation θ ,

$$L(r|\theta, a) = \frac{2}{N_0} \text{Re} \left\{ \int [r(t) - \frac{1}{2}s(t;\theta, a)]s^*(t;\theta, a) dt \right\}. \quad (3.40)$$

The received waveform $r(t)$ enters this expression through an inner product with the signal $s(t;\theta, a)$. Thus, in order to perform maximum-likelihood estimation on $\{\theta, a\}$, we must have available the deterministic reflection $s(t;\theta, a)$ for all orientations and target types. In practice, this is achieved through the development of databases, each of which contains the range-profiles for a given target at a discrete sampling of the possible orientations. The choice of a deterministic model then reduces to the choice of a particular algorithm for computing the range-profiles to be placed in the database. Recently, a number of efforts have been made toward the development of computer tools capable of simulating the data produced by a variety of remote sensors. One such tool that has been developed for radar simulations is known as XPATCH and is described in the next section.

HRR Simulator: XPATCH

The HRR signature prediction tool known as XPATCH is a computer software package, developed jointly by Wright Laboratory and DEMACO, Inc., that is designed to simulate the radar return from a given target [2]. The target is represented by a collection of triangular patches and outward normal vectors. Additionally, material parameters as well as edge and curvature information may be provided to further characterize the reflective properties of the target surface. The algorithm is based on a shooting and bouncing ray technique, in which rays are directed from the radar toward the target, and traced as they are reflected by the various surfaces on the

target. Those rays which are eventually directed back toward the radar contribute to the return. This technique allows for incorporation of the effects of both shadowing, in which portions of the target remain unilluminated because they lie in the shadow of other structures, and multi-bounce, in which the impinging wavefront may reflect many times inside cavities on the target prior to reflection toward the radar. When simulating the range-profile for a particular target, the algorithm is invoked to simulate the complex scattered far-field at discrete frequencies that are uniformly spaced within the specified transmitter bandwidth. The range-profile is found from the frequency-domain data through the Inverse Discrete Fourier Transform.

Given the choice of XPATCH as a deterministic model for the range-profile, it is desirable to know if the data simulated by this model support recognition, when processed using the deterministic log-likelihood function of (3.40). To investigate this possibility, a simple test was devised. The simulated data used for this test are drawn from a standardized collection of XPATCH range-profiles, collectively known as the University Research Initiative Synthetic Dataset. The following subsections describe the dataset, and go on to present the results of our test.

The URISD

The University Research Initiative Synthetic Dataset (URISD) is a collection of simulated range-profiles produced by the HRR simulator XPATCH. The URISD includes range-profiles simulated for four ground vehicles over three frequency bands and at three different elevation angles.

- Vehicles: 2 tank models, school bus, fire truck
- Frequency bands: UHF, L, X
- Elevation angles: 10° , 25° , 40°

Each vehicle is represented by a CAD model consisting of triangular patches covering the surface of the vehicle and a simulated rough ground surface, included so that the range-profiles exhibit the effects of reflections from the terrain in the vicinity of the vehicle. The tank models are generic in nature and are not intended to represent any particular existing tank; they are assigned the arbitrary labels of "m1" and "t1". The 36 data sets comprising the URISD each contain range-profiles for target azimuth angles that are uniformly sampled around the circle at a density that depends on the frequency band. Additionally, the data sets contain range-profiles simulated

Table 3.1: Frequency and azimuth parameters for the URISD.

Frequency Band	UHF	L	X
Center Frequency	600 MHz	1.5 GHz	10.0 GHz
Bandwidth	739.5 MHz	739.5 MHz	1.4775 GHz
Frequency Sample Spacing	4.921 MHz	4.921 MHz	4.921 MHz
Number of Frequency Samples	151	151	301
Azimuth Angle Spacing	0.3°	0.15°	0.02°
Number of Azimuth Samples	1201	2401	18001

for each of the four combinations of transmitted and received linear polarization states, $\{hh, hv, vh, vv\}$. The parameters for the various data sets are summarized in Figure 3.1.

We have made use of these data in developing a simulation of joint orientation tracking and recognition that is presented in Chapter 5. Additionally, these data were used in a number of simple tests presented in the remainder of this chapter to compare the recognition performance associated with candidate HRR models and the associated likelihood functions.

Recognition Test: L-band Data

Range-profiles were drawn from the URISD for the purpose of testing the recognition performance under the deterministic log-likelihood function (3.40). For the first test, we selected a subset of the L-band data; the range-profiles for each of the four targets at an elevation of 10 degrees and for azimuth angles over the range of 10 to 30 degrees were extracted and placed into libraries for each target. Also for this test, only the data for the vv polarization state were used. The range-profiles are depicted in Figure 3.2. For ease in viewing, the figure shows only those range-profiles for azimuths between 10 and 15 degrees. The upper left panel shows the range-profiles for a fire truck. The lower left panel shows the range-profiles for a school bus. The right panels show the range-profiles for the two tank models.

As evidenced by the figure, the range-profiles produced by XPATCH exhibit extreme variability for small changes in azimuth angle. This behavior is in general agreement with that observed for real data. To test whether this variability allows for the use of these data for recognition, the range-profile from the fire truck dataset at an azimuth angle of 19.95 degrees was selected as an observation, and the log-likelihood

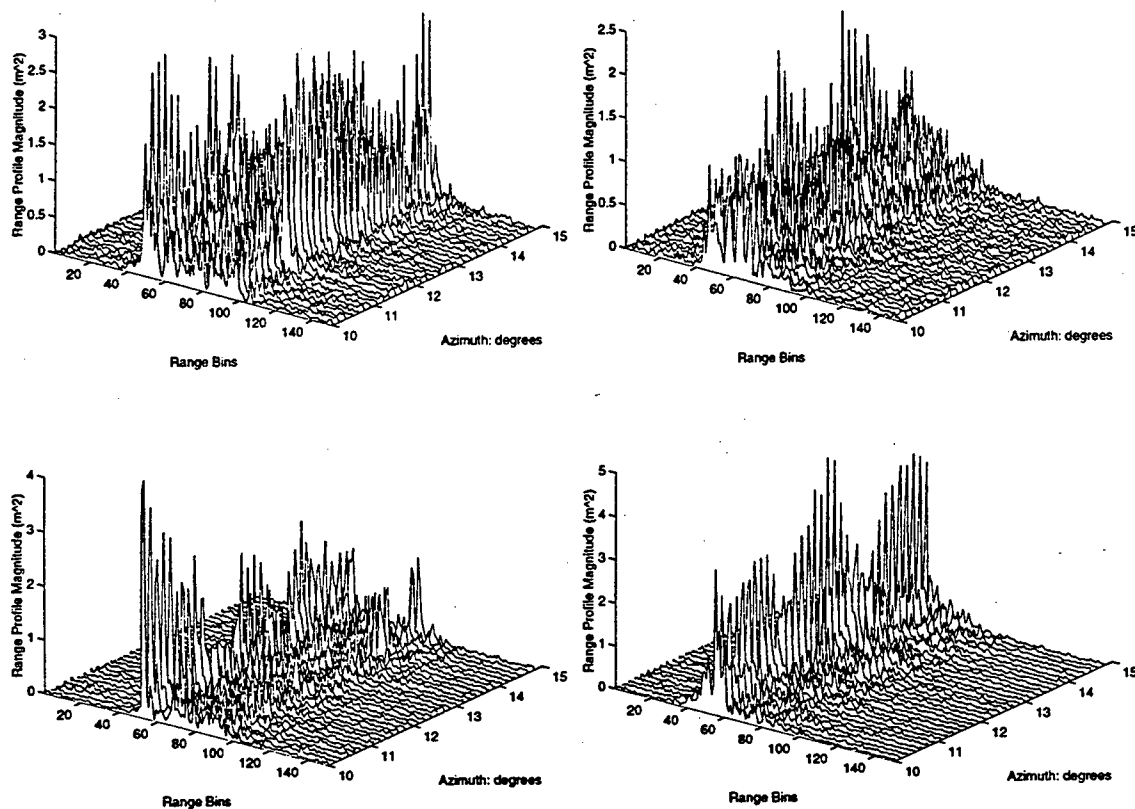


Figure 3.2: Simulated L-band range-profiles for four ground vehicles over a 5 degree sweep in azimuth. The upper left panel shows the range-profiles for a fire truck. The lower left panel shows the range-profiles for a school bus. The right panels show the range-profiles for two tank models.




Figure 3.3: Deterministic model log-likelihood function for a four-vehicle recognition test using L-band data. The range-profile for the fire truck at 19.95 degree azimuth was selected as an observation, and the deterministic model log-likelihood was computed for each vehicle for azimuth angles from 10 to 30 degrees. The right panel is an enlargement of the left panel showing the behavior in the vicinity of the true azimuth. In both panels, the solid line shows the log-likelihood for the correct vehicle, the fire truck.

function in equation (3.40) was computed for each of the four vehicles over the range of azimuth angles from 10 to 30 degrees. The results of this test are depicted in Figure 3.3. The left panel of the figure shows the log-likelihood for all azimuth angles, while the right panel is an enlargement showing the behavior near the true azimuth angle. In both panels, the solid line shows the log-likelihood for the correct vehicle, the fire truck. The results of this test show that recognition is possible using these data, in that the likelihood of the correct vehicle in the vicinity of the true azimuth angle is larger than the likelihood for any of the incorrect vehicles. However, the figure also shows that the azimuth angle must be estimated to within ± 0.5 degrees of the truth for recognition to be successful. These results are in accordance with previous simulations using other targets.

It is important to note that no noise was added to the observation prior to conducting this test. However, we cannot set $N_0 = 0$ for the likelihood calculation, as N_0 appears in the denominator of (3.40). We therefore chose to set $N_0 = 0.01$ square meters for the computation of the likelihood function depicted in Figure 3.3. This value was chosen both because it represents a high signal-to-noise ratio, and for comparison with some results appearing later in this chapter in which the same noise level is used with a different model for the data. We may therefore interpret

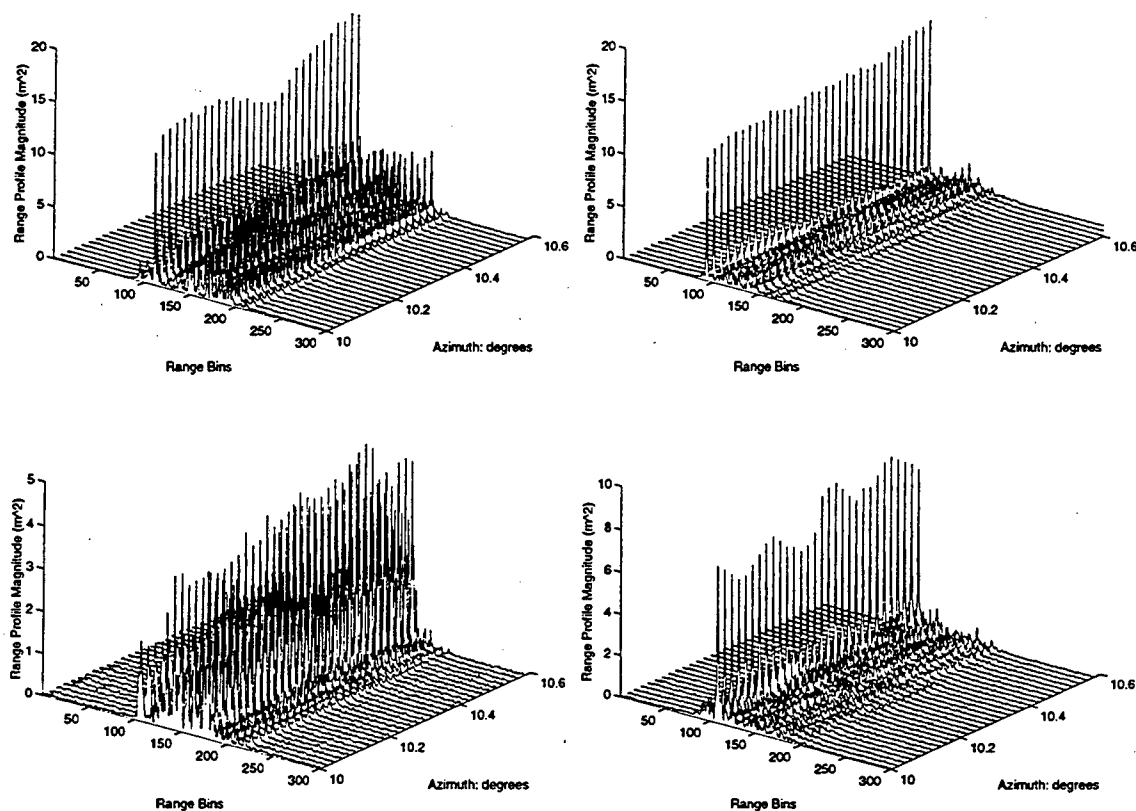


Figure 3.4: Simulated X-band range-profiles for four ground vehicles over a 0.6 degree sweep in azimuth. The upper left panel shows the range-profiles for a fire truck. The lower left panel shows the range-profiles for a school bus. The right panels show the range-profiles for two tank models.

the plots of Figure 3.3 as depicting the expected value of the log-likelihood function, where the expectation is over the ensemble of realizations of the noise process. Given this interpretation of the figure, we can now say that for this case the log-likelihood of the correct vehicle in the vicinity of the true azimuth angle is larger than that of any of the incorrect vehicles by a margin of more than 1500.

Recognition Test: X-band Data

The test presented above was repeated using X-band data from the URISD. As before, the data for each of the four vehicles at 10 degrees elevation, for vv polarization, and for azimuth angles between 10 and 30 degrees were placed in libraries. Selected range-profiles from these libraries are shown in Figure 3.4. The range-profile for the fire




Figure 3.5: Deterministic model log-likelihood function for a four-vehicle recognition test using X-band data. The range-profile for the fire truck at 19.95 degree azimuth was selected as an observation, and the deterministic model log-likelihood was computed for each vehicle for azimuth angles from 10 to 30 degrees. The right panel is an enlargement of the left panel showing the behavior in the vicinity of the true azimuth. In both panels, the solid line shows the log-likelihood for the correct vehicle, the fire truck.

truck at 19.95 degrees azimuth was selected as an observation, and the deterministic log-likelihood function was computed for each vehicle and azimuth. The results of this test are presented in Figure 3.5.

Again the figure shows a positive result for recognition, as the likelihood for the correct target in the vicinity of the true azimuth angle is larger than that for any other vehicle, but the azimuth angle must be known to within 0.1 degrees of the truth for correct recognition. It is unknown if such accuracy in orientation estimation is achievable in any realistic setting. As in the L-band test, we have set $N_0 = 0.01$ in computing the log-likelihood function, so that the plot in Figure 3.5 depicts the expected value of the log-likelihood function over the ensemble of noise realizations. We see that the log-likelihood for the correct target at the true azimuth is larger than that for any other target by a margin of 10^4 .

Limitations of the Deterministic Model

The practical problem with the deterministic model as demonstrated in these tests is that all of the variability of range-profiles with changing target orientation is reflected in the likelihood function. It would be far more desirable if the model for range-profiles captured those features of the data that are relatively invariant over small changes in

orientation. It has been the approach of some investigators to choose a particular feature, over that it is invariant to target orientation, and to perform recognition directly in the feature space. This is not an entirely unreasonable approach, as examination of simulated and real range-profiles reveals that certain observably invariant features seem to exist. For example, it has been widely reported that for small changes in orientation, the amplitudes of peaks in the range-profile change drastically, while the range locations of the dominant peaks do not. Rather than choosing a particular set of features for recognition, in the next section we consider a simple stochastic model that characterizes the statistics for range-profiles for a particular target over a small region in orientation.

The fundamental assumption inherent in the use of a deterministic model is that the range-profile is computable given the target orientation. This may or may not be physically accurate. For aircraft targets, even ignoring propagation effects, the airframe flexes during flight introducing variations in the range-profiles by moving relative distances of pairs of points through more than a wavelength. Similarly, ground vehicles experience vibration induced by the propulsion system. This flexing or vibration itself may be enough to invalidate the deterministic model. Other physical effects such as moisture or dirt on the target and control surface movement for airframes or wheel rotation for ground vehicles are not adequately modeled by signature prediction tools. These indicate that stochastic models (based in large part on the outputs of tools such as XPATCH) may be best for HRR range-profiles. In the next section, we introduce a candidate stochastic model.

3.2.2 Conditionally Gaussian Model

In this section we consider a stochastic radar model under which, given the target type and orientation, the signal portion $s(t; \theta, a)$ of the received range-profile forms a complex Gaussian random process. The conditional mean and covariance are denoted as $\mu(t|\theta, a)$ and $K(t_1, t_2|\theta, a)$, respectively. The WSSUS model assumes that returns from different delays τ are statistically uncorrelated. For our conditionally Gaussian model, we extend this assumption to samples of the range-profile.

Let the model be discretized such that samples of the received range-profile $r(t)$ are collected in the column vector \mathbf{r} . Sampling is performed by integrating the received range-profile against unit-energy signals, typically rectangular pulses, of

duration equal to the sampling interval, as discussed earlier. Then the model is

$$\mathbf{r} = \mathbf{s}(\theta, a) + \mathbf{w}, \quad (3.41)$$

where, given $\{\theta, a\}$, \mathbf{s} is a Gaussian random vector with mean $\mu(\theta, a)$ and covariance matrix $\mathbf{K}(\theta, a)$, and \mathbf{w} is also Gaussian with mean 0 and covariance $N_0\mathbf{I}$. Thus,

$$E\{\mathbf{r}|\theta, a\} = \mu(\theta, a) \quad (3.42)$$

$$E\{[\mathbf{r} - \mu(\theta, a)]^\dagger [\mathbf{r} - \mu(\theta, a)]|\theta, a\} = \mathbf{K}(\theta, a) + N_0\mathbf{I}, \quad (3.43)$$

where † denotes complex conjugate transposition. Under this model, the log-likelihood of observing \mathbf{r} given $\{\theta, a\}$ is

$$l(\mathbf{r}|\theta, a) = -\ln |\mathbf{K} + N_0\mathbf{I}| - [\mathbf{r} - \mu]^\dagger (\mathbf{K} + N_0\mathbf{I})^{-1} [\mathbf{r} - \mu], \quad (3.44)$$

where, for notational simplicity, we have suppressed the explicit dependence of \mathbf{K} and μ on θ and a . *The key assumption of the conditionally Gaussian model is that the returns from different range bins are statistically uncorrelated, so that \mathbf{K} is a diagonal matrix.* There are several possible assumptions on $s(t; \theta, a)$ that would yield independence of the range samples. One is that the sampling is performed by integrating $\mathbf{r}(t)$ against eigenfunctions of the autocorrelation function for $s(t; \theta, a)$. Another is that $s(t; \theta, a)$ is uncorrelated with itself at different times t ,

$$K(t_1, t_2|\theta, a) = K(t_1|\theta, a)\delta(t_1 - t_2), \quad (3.45)$$

and that sampling uses non-overlapping rectangular pulses.

Recognition Tests: L-band and X-band

We implemented this model on the simulated HRR datasets used in the previous section for the purpose of testing recognition performance. A sliding window in orientation was defined. The size of the window was arbitrarily chosen to be 3 degrees for the L-band data and 1 degree for the X-band data. Initially, the edge of the window was set to the minimum azimuth for all the range-profiles in the dataset, namely 10 degrees. All of the range-profiles corresponding to azimuth angles inside the window were drawn from the dataset for a particular vehicle, and the mean range-profile and the variance within each range bin were computed. The window was shifted by one




Figure 3.6: Conditionally Gaussian model log-likelihood functions for four-vehicle recognition tests. The range-profile for the fire truck at 19.95 degree azimuth was selected as an observation, and the conditionally Gaussian log-likelihood was computed for each vehicle for azimuth angles from 10 to 30 degrees. The left panel shows the results for L-band data, while the right panel shows the results for X-band data. In both panels, the solid line shows the log-likelihood for the correct vehicle, the fire truck.

position, and the computation was repeated. In this manner, the mean μ and the diagonal elements of \mathbf{K} were determined for each window position over the azimuth angles from 10 to 30 degrees. This process was repeated for each of the four vehicles.

A range-profile for a particular vehicle at a particular azimuth was selected as an observation \mathbf{r} , and the log-likelihood (3.45) was computed for each position of the sliding window and each vehicle. Because no noise was added to the observation, N_0 was set to zero for the log-likelihood computation.

The results of this test are presented in Figure 3.6. The observation was selected as the range-profile for the fire truck at 19.95 degrees azimuth. It is clear from the figure that a very significant improvement in performance was achieved for this example. The log-likelihood for the correct vehicle is larger than that for any other vehicle by a margin of more than 200 for the L-band data, and by a margin of more than 3000 for the X-band data. More importantly, this log-likelihood margin is stable over the entire range of azimuth angles considered, so that correct recognition is possible for this example even when the azimuth is in error by 10 degrees.

Range Windowing

This test was modified and repeated in order to validate the results shown above. The modification was to apply a window to all of the range-profiles in the dataset for each target, such that the returns from range bins corresponding to the simulated rough ground were eliminated, and only returns from the range extent of the vehicle remained. The determination of which range bins to select was made by visual inspection of the data, and the same range bins were selected from range-profiles for all vehicles at all azimuth angles. This modification was performed to eliminate the influence of the returns from the terrain on recognition. It has been observed that the variance of the returns from these range bins is very small, thereby possibly over-weighting these returns in the log-likelihood calculation.

Given this modification, the test was repeated using X-band data and for the same observation used previously, the fire truck at 19.95 degrees azimuth. The modified test was also applied for observations of other vehicles at other azimuth angles. The results of these tests are shown in Figure 3.7. The upper left panel shows the results when observing the fire truck for comparison with the right panel of Figure 3.6. Such a comparison reveals that the elimination of range bins corresponding to the simulated terrain resulted in a slight lowering of the log-likelihood for all targets at all angles by a constant amount, but did not affect recognition performance in any significant way. The lower left panel shows the result when the observation is the school bus at 25 degrees azimuth. The two right panels show the results when the observations are each of the two tank models at 15 degrees azimuth. In each panel, the solid line shows the log-likelihood for the fire truck, the dotted line represents the m1 tank, the dashed line represents the school bus, and the dot-dash line represents the t1 tank.

The results of these tests are somewhat less positive than those of Figure 3.6. Specifically, for the observation of the school bus at 25 degrees azimuth or the t1 tank at 15 degrees azimuth, the log-likelihood margin for the correct vehicle for azimuth angles in the vicinity of the truth is approximately 300, and significant errors in azimuth angle can lead to incorrect recognition. For the school bus, the azimuth angle must be known to within ± 2 degrees of the truth for successful recognition.

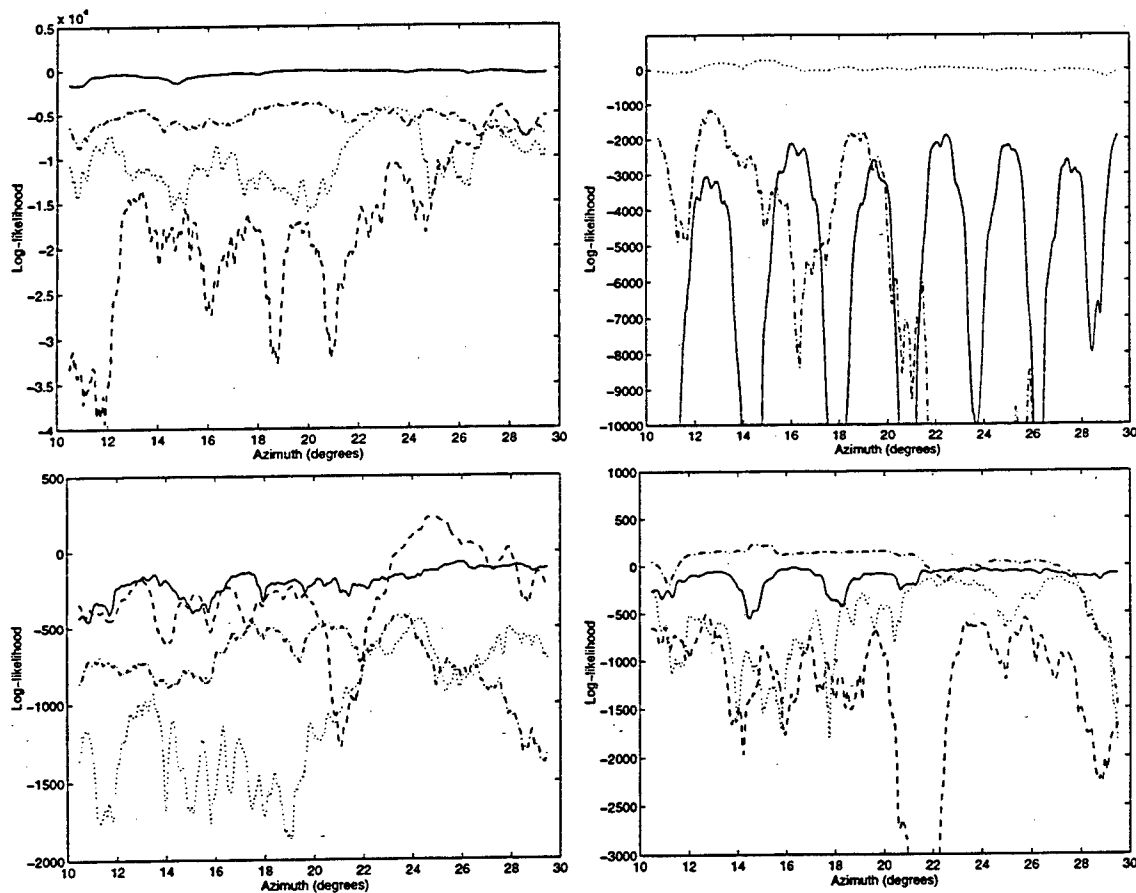


Figure 3.7: Conditionally Gaussian model log-likelihood functions for four-vehicle recognition tests. In the upper left panel, the range-profile for the fire truck at 19.95 degrees azimuth was selected as an observation, and the conditionally Gaussian log-likelihood was computed for each vehicle for azimuth angles from 10 to 30 degrees. The lower left panel shows the result when the observation is the school bus at 25 degrees azimuth. The two right panels show the results when the observations are each of the two tank models at 15 degrees azimuth. In each panel, the solid line shows the log-likelihood for the fire truck, the dotted line represents the m1 tank, the dashed line represents the school bus, and the dot-dash line represents the t1 tank.

Artificial Noise Variance

A further modification of the tests shown above was conducted. The modification was to introduce a non-zero value for N_0 to the conditionally Gaussian log-likelihood function (3.44), even though the observation contains no additive noise. The concern is that if the computed variance in some range bins is very small, then the returns from these range bins will be overemphasized in the log-likelihood function. By adding a small constant amount to the diagonal of the covariance matrix, the returns from such range bins will be de-emphasized, possibly resulting in improved performance. The results of this modification for a value of N_0 of 0.01 square meters are summarized in Figure 3.8. In each panel, the solid line represents the log-likelihood for the correct vehicle, and the dot-dash line represent the log-likelihood for the next most likely vehicle, each for a noise variance of 0, as shown in the corresponding panel of Figure 3.6. The dotted line represents the log-likelihood for the correct target, and the dashed line represents the log-likelihood for the next most likely target, each for a noise variance of 0.01.

This test allows for a direct numerical comparison of the log-likelihood values under our two candidate target models. The upper left panel of Figure 3.8 indicates that the log-likelihood for the correct vehicle at the true azimuth is larger than that for any other vehicle by a margin of roughly 3000. In the previous subsection, we saw that under the deterministic model, using the same correct vehicle, true azimuth, and noise variance, the log-likelihood margin was 10^4 . Thus, this example shows that the log-likelihood margin is actually larger under the deterministic model, but this margin is achieved over an exceedingly narrow range in azimuth.

The plots of Figure 3.8 indicate that the recognition performance was not significantly altered by the addition of a small constant value to the variance. The exception to this was the case where the school bus was observed, shown in the lower left panel of Figure 3.8. Notice that the deep troughs in the likelihood function for the correct target were shallowed by the modification. This suggested that the recognition performance under the conditionally Gaussian model might be improved by this technique for the proper choice of N_0 . The test was repeated using a value of N_0 of 0.1 square meters. The results of this test are shown in Figure 3.9. In each panel, the solid line represents the log-likelihood for the correct vehicle, and the dot-dash line represent the log-likelihood for the next most likely vehicle, each for a noise variance of 0, as shown in the corresponding panel of Figure 3.6. The dotted

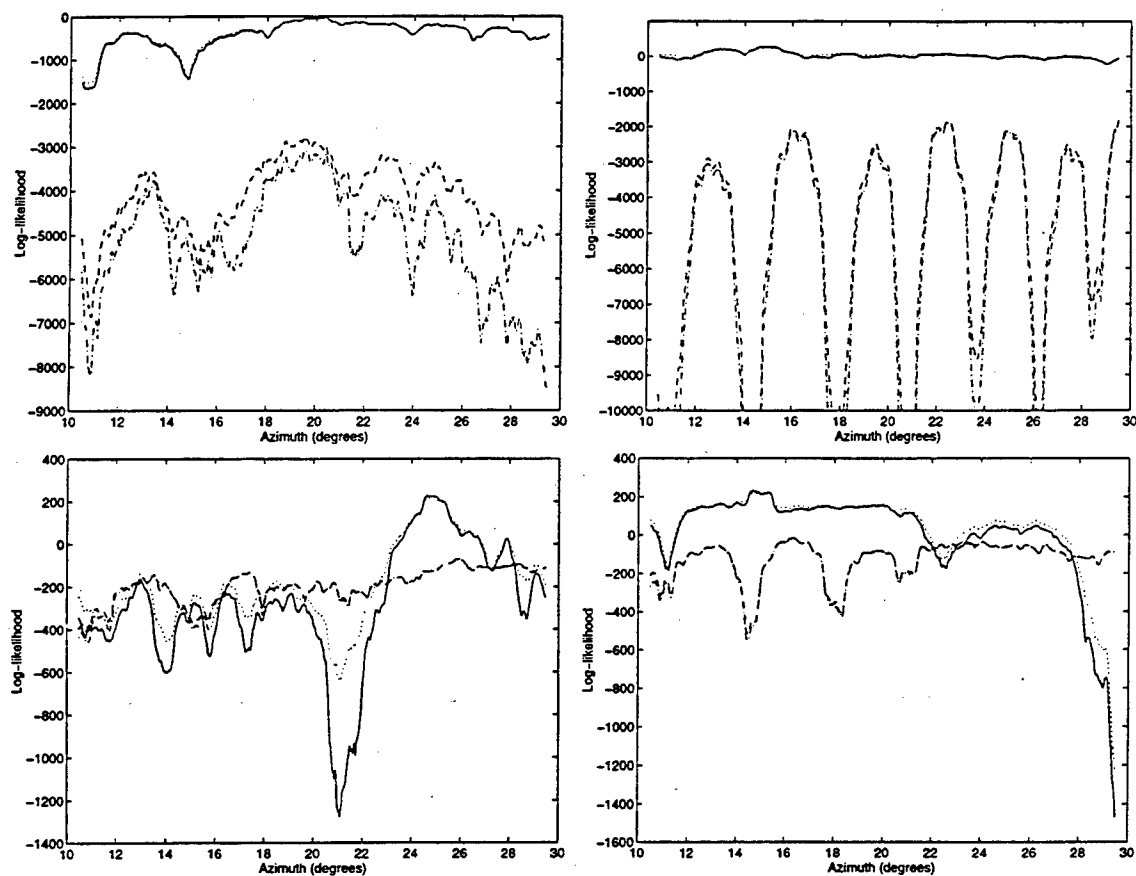


Figure 3.8: Comparison of conditional Gaussian log-likelihood functions for noise variances of 0 and 0.01 square meters. In each panel, the solid line represents the log-likelihood for the correct vehicle, and the dot-dash line represent the log-likelihood for the next most likely vehicle, each for a noise variance of 0, as shown in the corresponding panel of Figure 3.6. The dotted line represents the log-likelihood for the correct target, and the dashed line represents the log-likelihood for the next most likely target, each for a noise variance of 0.01.

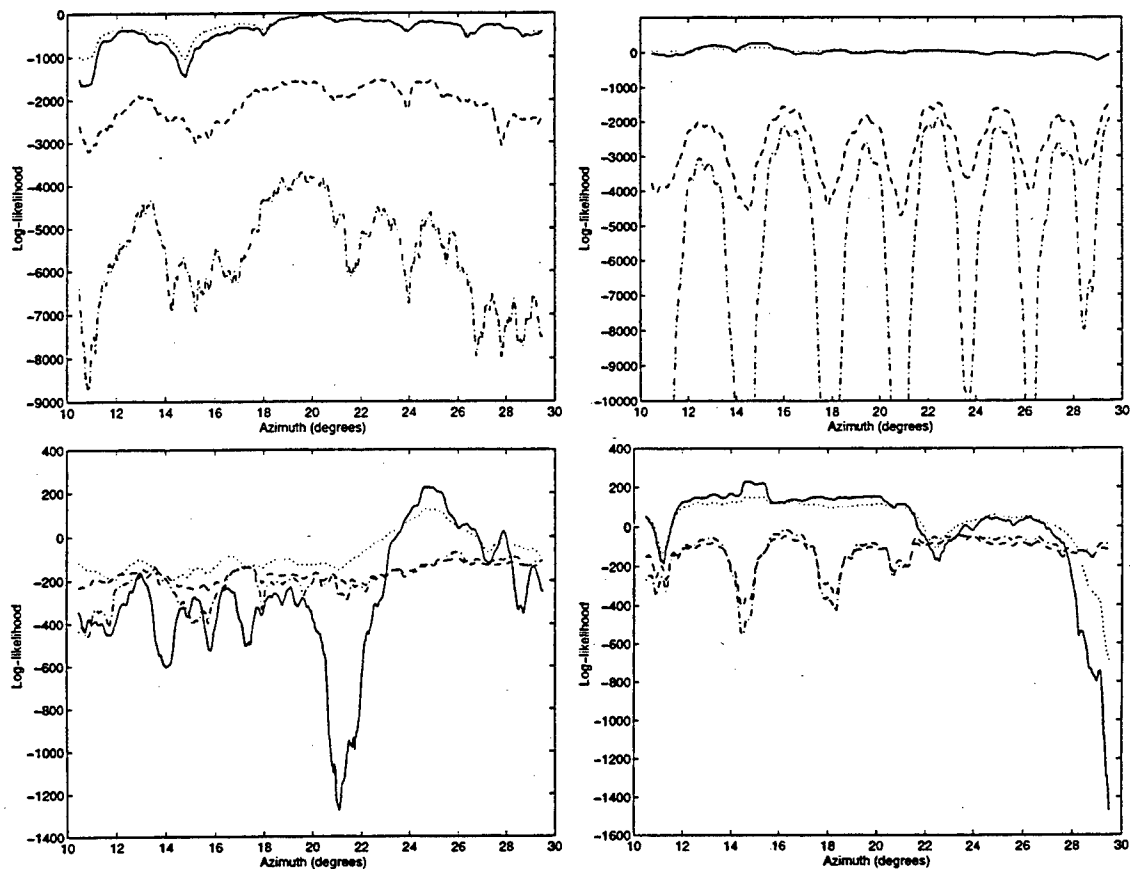


Figure 3.9: Comparison of conditional Gaussian log-likelihood functions for noise variances of 0 and 0.1 square meters. In each panel, the solid line represents the log-likelihood for the correct vehicle, and the dot-dash line represent the log-likelihood for the next most likely vehicle, each for a noise variance of 0, as shown in the corresponding panel of Figure 3.6. The dotted line represents the log-likelihood for the correct target, and the dashed line represents the log-likelihood for the next most likely target, each for a noise variance of 0.1.

line represents the log-likelihood for the correct target, and the dashed line represents the log-likelihood for the next most likely target, each for a noise variance of 0.1.

Inspection of the plots of Figure 3.9 might initially suggest that the recognition performance was worse for the larger value of $N_0 = 0.1$ square meters, as the log-likelihood for the incorrect target is in each case increased after the modification. However, it is still the case that correct recognition is achievable for all for cases, albeit with a somewhat decreased log-likelihood margin. A very interesting result occurred for the observation of the school bus, shown in the lower left panel. In this case, the troughs in the log-likelihood were almost completely eliminated, with the result that correct recognition is achievable over a much wider range of azimuth angles as a result of this modification. Thus, all four cases represented by Figure 3.9 indicate that correct recognition is achievable under the modification of the conditionally Gaussian model, even when the azimuth angle is in error by 5 degrees or more.

Recognition Test for Broadside Aspect

As a further comparison between the recognition performance under the deterministic and conditionally Gaussian models, a search was conducted for instances in which the performance under the conditionally Gaussian model was particularly poor. Such a case was found when the school bus was observed at azimuth angles near 90 degrees or broadside to the radar. The recognition problem for this case is expected to be difficult, as all four vehicles possess large, flat lateral surfaces that are normal to the radar as these azimuths. The range-profiles for all four vehicles at these azimuths are dominated by a strong peak at the range location of the lateral surface. This similarity in the data is shown in Figure 3.10. The range-profile for the school bus at 90 degrees azimuth was selected as an observation and the log-likelihoods under both models were computed for all four vehicles over the range of azimuth angles from 80 to 100 degrees. The resulting log-likelihood functions are shown in Figure 3.11. The left panel shows the log-likelihood for the deterministic model. The right panel shows the log-likelihood for the conditionally Gaussian model. Both panels in this figure are enlargements that show the behavior near the true azimuth angle. The results of this test show that recognition is indeed difficult for this case. The deterministic model results in a log-likelihood margin of more than 6×10^5 , but only if the azimuth angle is known to within 0.1 degrees of the truth. The conditionally Gaussian model results in a log-likelihood margin of more than 1000, but only if the azimuth angle is known to within 0.5 degrees of the truth. For this example, use of the conditionally

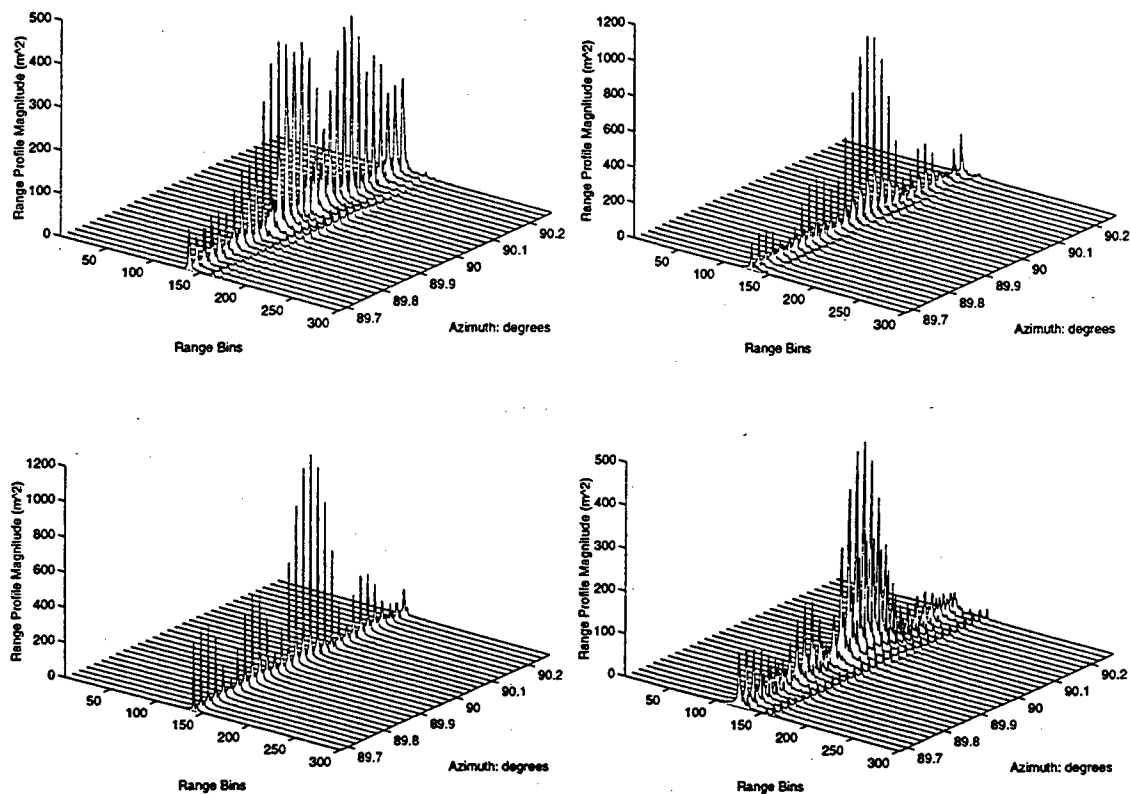


Figure 3.10: Simulated X-band range-profiles for four ground vehicles near 90 degrees azimuth. The upper left panel shows the range-profiles for a fire truck. The lower left panel shows the range-profiles for a school bus. The right panels show the range-profiles for two tank models.




Figure 3.11: Log-likelihood functions for four-vehicle recognition tests near 90 degrees azimuth. The observation is the range-profile for the school bus at 90 degrees azimuth. The left panel shows the log-likelihood for the deterministic model. The right panel shows the log-likelihood for the conditionally Gaussian model. In both panels, the solid line is for the correct vehicle, the school bus.

Gaussian model resulted in an increase in the range of azimuth angles over which correct recognition is achievable, but also resulted in a decreased log-likelihood margin for the correct vehicle. It is interesting to note that the width of the range in azimuth angles over which the log-likelihood for the conditionally Gaussian model is approximately constant is roughly equal to the width of the orientation window over which the mean and variance were computed. This suggests that the recognition performance may be stabilized over larger ranges in azimuth by increasing the width of the sliding window.

Recognition Test: NVESD Millimeter-Wave Radar

We have conducted an additional recognition test using simulated millimeter-wave radar data. We obtained CAD models for two ground vehicles, a tank and a HMMWV personnel vehicle, and used XPATCH to simulate range-profiles for these vehicles. For these simulations the elevation angle was held constant at one degree, and the azimuth angle was varied between 10 and 30 degrees in increments of 0.01 degrees.

The frequency parameters for the XPATCH simulations were motivated by real millimeter-wave range-profiles obtained from the U.S. Army Night Vision Electronic Sensors Directorate. Examples of these range-profiles are shown in Figure 3.12. Examination of the range-profiles reveals that there are approximately 20 range bins




Figure 3.12: Real millimeter-wave range-profiles of a ground vehicle of unknown type. The left panel shows the range-profile for horizontal polarization during transmission and reception, and the right panel shows the corresponding range-profile for vertical polarization during transmission and reception. These data were provided courtesy of the U.S. Army Night Vision Electronic Sensors Directorate.

in which the magnitude is larger than the background. If this region defines the target extent in the range-profile, and if the target extent in the range dimension is 10 meters, then the size of each range bin is 0.5 meters. Furthermore, if B denotes the bandwidth, and if the return is sampled at a rate of $\frac{1}{B}$ so that the size of a range bin is equal to the range resolution, then B is approximately 300MHz. Given that there are 64 frequency samples in the data provided to us, the frequency spacing between samples should be somewhat less than 5MHz. We were informed by NVESD that the center frequency for this radar is approximately 35GHz. Therefore, we chose the following frequency parameters for our simulations. The center frequency was set to 35GHz, and the return was simulated at 64 discrete frequencies with 5MHz between adjacent frequencies. The resulting range resolution of the simulated data is 0.47 meters.

Examples of the simulated range-profiles are shown in Figure 3.13. The left panel shows the range profiles for the HMMWV at a fixed elevation angle of 1 degree and for azimuth angles between 10 and 10.3 degrees. The right panel shows the range profiles for the tank model at the same set of orientations.

We performed a recognition test similar to those presented earlier using the simulated millimeter-wave range-profiles. The range-profile for the tank at 20 degrees azimuth was selected as an observation, and the log-likelihood for both vehicles at




Figure 3.13: Simulated millimeter-wave range profiles for two ground vehicles. The left panel shows the range profiles for a HMMWV personnel vehicle at a fixed elevation angle of 1 degree and for azimuth angles between 10 and 10.3 degrees. The right panel shows the range profiles for a tank model at the same set of orientations.

azimuth angles between 10 and 30 degrees was computed. The resulting log-likelihood functions under the two models are shown in Figure 3.14. This figure displays many of the features seen in our previous recognition tests. Recognition is possible under both models in that the likelihood of the correct target at the true orientation is larger than the likelihood for the incorrect target. The likelihood margin between the correct target and incorrect target at the true azimuth is larger under the deterministic model than under the Gaussian model. However, the deterministic model is not robust with respect to orientation as there are many orientations at which the likelihood of the incorrect target is larger than that of the correct target. In contrast, the Gaussian model is robust with respect to orientation.

Model Comparison Using Rome Laboratory Real Range-Profiles

The various investigations of our two candidate models for HRR data have thus far made exclusive use of simulated range-profiles. In this section we extend the comparison to include real data. We have been provided with HRR range-profiles collected at the U.S. Air Force Rome Laboratory radar facility. In this experiment, targets of opportunity were illuminated by an S-band radar. The transmitted signal was a linear FM chirp, with a center frequency of 3.408 GHz and a bandwidth of 320MHz. Vertical polarization was used during transmission and reception. Each data set consists of 119 range-profiles collected over a time interval of less than one




Figure 3.14: Log-likelihood functions for a two-vehicle recognition test using simulated millimeter-wave range-profiles. The left panel shows the log-likelihood under the deterministic model and the right panel shows the log-likelihood under the Gaussian model. In both panels, the solid line represents the result for the correct vehicle, the tank.

second, with successive range-profiles occurring approximately 4 milliseconds apart. Figure 3.15 displays the range-profiles from three of the aircraft collected in this experiment. The left panel shows the range-profiles for a DC-10 aircraft. The center and right panels show the range-profiles for DC-9 and Cessna 310 aircraft, respectively.

We wish to use these data to perform a model comparison testing the potential for recognition under each model, similar to those conducted previously with simulated data. However, due to several characteristics distinguishing these real range-profiles from the simulated data, we have made significant modifications to the computations used for model comparison. First of all, we have available to us only the real data here, including receiver noise and any other distortions resulting from the collection process. As a result, these noisy observations will be used in place of both the observed data r and the signal portion s in our likelihood calculations. Also, the change in the orientation of the illuminated aircraft with respect to the radar system between adjacent range profiles is exceedingly small. Therefore, we have used all 119 range-profiles from each aircraft to compute a single template under each model. For the Gaussian model, the template will consist of the mean range-profile and the variance in each range-bin as before. For the deterministic model, the template will consist of the mean range-profile. This is a significant change from the computations performed using simulated data, in which the deterministic model templates consisted of individual range-profiles drawn from the dataset. Also, for




Figure 3.15: Real S-band range-profiles for three commercial aircraft. Each panel displays 119 range profiles collected at 4 millisecond intervals. The left panel shows the range-profiles for a DC-10 aircraft. The center and right panels show the range-profiles for DC-9 and Cessna 310 aircraft, respectively. Data courtesy of the U.S. Air Force Rome Laboratory.

both models, we will use each individual range-profile from each aircraft as a distinct observation to compare to the single template. This is another difference from the case of simulated data, in which we drew a single range profile from the dataset and declared it to be the observation. Additionally, we do not know the noise variance associated with these data. We have therefore chosen an arbitrary value for N_0 of 0.1 for the likelihood computations. Note that N_0 is not in units of square meters here, as the range-profiles are not calibrated in these units. Finally, we do not know the aircraft orientation at the illumination times, and it is likely that the orientations assumed by the three aircraft during the experiment are completely different.

With all of the listed caveats, Figure 3.16 shows the results of our model comparison. The left panel shows the resulting log-likelihood functions under the deterministic model, and the right panel shows the log-likelihoods under the Gaussian model. In both panels, the solid line represents the data for the correct aircraft, the DC-10, the dashed line represents the DC-9, and the dotted line represents the Cessna 310. While there were many differences between this model comparison and those presented previously using simulated data, the results depicted in the figure are in general agreement with the conclusions we have already drawn with respect to the models. The log-likelihood margin between the correct target and the incorrect target is, at many orientations, larger under the deterministic model. However, the deterministic model is not robust with respect to small changes in orientation, and there are several orientations at which incorrect recognition would result under the deterministic model. In contrast, the Gaussian model is robust with respect to small changes in orientation, at the expense of a lowered log-likelihood margin. There are




Figure 3.16: Model comparison using real HRR range-profiles. Each of the 119 range-profiles for the DC-10 aircraft were compared to a single template for each aircraft. The left panel shows the resulting log-likelihood functions under the deterministic model, and the right panel shows the log-likelihoods under the Gaussian model. In both panels, the solid line represents the data for the correct aircraft, the DC-10, the dashed line represents the DC-9, and the dotted line represents the Cessna 310.

no orientations at which incorrect recognition would occur under the conditionally Gaussian model.

Clearly this is only a preliminary comparison of our two models using real data. It is desired to perform a more extensive comparison, using many more range-profiles collected for more targets, and where the target orientation at the illumination times is known to within some specified tolerance. We are currently working with personnel from the U.S. Army Missile Command to obtain data for such a test.

Recognition Test for Variation in Two Orientation Angles

The tests presented above have investigated the recognition performance under our two models for small changes in a single orientation angle. To extend the results of these tests, we have conducted similar tests using data from aircraft targets in which two orientation angles are varied. For this test, CAD models for the X-29 and the F-117 aircraft were obtained, and XPATCH was used to simulate range-profiles for these targets at a discrete set of orientations defined by the yaw and pitch angles of the target relative to the sensor. Both the yaw and pitch angles were varied over the range from -10 to 10 degrees, where 0 degrees yaw and 0 degrees pitch is defined to be the nose-on orientation. The yaw angle was sampled uniformly at 0.2 degree




Figure 3.17: Simulated S-band range-profiles for two aircraft models. The left panel shows the range-profiles for the F-117 aircraft at a fixed yaw angle of -0.4 degrees and for pitch angles in between -10 and -4 degrees. The right panel shows the range-profiles for the X-29 aircraft at the same set of orientations.

intervals, and the pitch angle was sampled at a yaw-dependent interval chosen to keep the arc length on the unit sphere between neighboring pitches approximately 0.2 degrees. This sampling scheme resulted in a discrete set of $10,165$ orientations for each aircraft.

Examples of the range-profiles for each of the two aircraft are shown in Figure 3.17. The left panel shows the range-profiles for the F-117 aircraft at a fixed yaw angle of -0.4 degrees and for pitch angles in between -10 and -4 degrees. The right panel shows the range-profiles for the X-29 aircraft at the same set of orientations. The frequency parameters for these simulations were chosen to approximate those of the S-band radar facility at Rome Laboratory described in the previous section: the center frequency is 3.2GHz , the bandwidth is 320MHz , and the range resolution is 0.46 meters.

The range-profiles produced by XPATCH were collected into libraries for use in the recognition tests. For the deterministic model, all $10,165$ range-profiles for each aircraft were placed in the libraries. For the conditionally Gaussian model, a moving patch in orientation was defined. The size of the patch was 1.2 degrees in both yaw and pitch. To begin, the edges of the patch in both yaw and pitch were set to -10 degrees. All range-profiles corresponding to orientations within the patch were collected, and the mean range-profile and the variance in each range bin were computed. The patch was shifted by 0.6 degrees in pitch, and the mean and variance were again computed. This process was repeated, each time shifting the patch by 0.6

degrees in yaw or pitch, to produce 961 overlapping patches for yaw and pitch between -10 and 10 degrees. The libraries for the conditionally Gaussian model contain the mean and variance range-profiles computed from each patch.

Novel range-profiles were simulated for each of the aircraft at a yaw angle of -4.34 degrees, and at a pitch angle of 4.37 degrees. These range-profiles were used as the observations for the recognition tests. Thus, for the tests that follow, the observed range-profile appears nowhere in the libraries for the deterministic model, and is not used in computing any library element for the stochastic model.

For the first test, the observed range-profile for the F-117 was selected, and the deterministic model log-likelihood was computed for both aircraft at all orientations. Although no noise was added to the observation used for this test, the log-likelihood was computed using a value for N_0 of 0.1 square meters. The results of this test are shown in Figure 3.18. The upper left panel shows the log-likelihood surface for the F-117 aircraft at varying orientations, and the upper right panel shows the log-likelihood surface for the X-29 aircraft. To aid in interpreting these results, the lower panels show the difference between the log-likelihoods for the two aircraft as a function of yaw and pitch. The lower right panel shows this difference as an image. Regions where the log-likelihood for the correct aircraft is higher than that for the incorrect aircraft are shown in white, regions where the log-likelihood for the incorrect aircraft is higher are shown in gray, and the true orientation is marked by an "X". From this panel it is clear that even a small error in orientation may lead to incorrect recognition for this case.

This test was repeated under the conditionally Gaussian model. The log-likelihood of the F-117 observation arising from each of the 961 orientation patches was computed for each aircraft type. To enable comparison with the results under the deterministic model, N_0 was set to 0.1 square meters. The results of this test are shown in Figure 3.19. The upper left panel shows the log-likelihood surface for the F-117 aircraft at varying orientations. The upper right panel shows the log-likelihood surface for the X-29 aircraft. The lower panel shows the difference between the log-likelihoods for the two aircraft as a function of yaw and pitch. Note that the log-likelihood difference is positive for all orientations. Therefore, an image is not shown in this figure, as it would be everywhere white. The results of this test indicate better performance for the conditionally Gaussian model, in the sense that correct recognition is possible when the orientation estimate is in error by several degrees. Under the deterministic model, small errors in the orientation estimate can

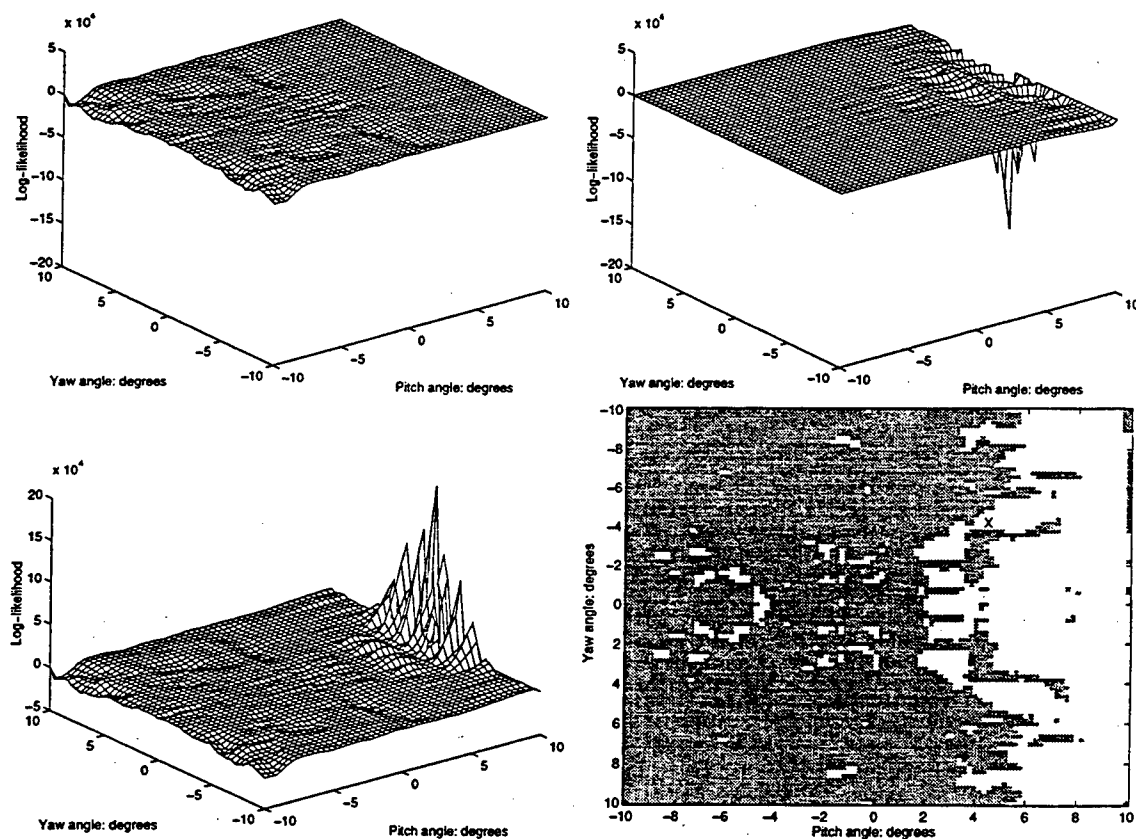


Figure 3.18: Deterministic model log-likelihood functions for a two-aircraft recognition test. The observation was the range-profile for the F-117 aircraft at -4.34 degrees yaw and 4.37 degrees pitch. The upper left panel shows the log-likelihood surface for the F-117 aircraft at varying orientations. The upper right panel shows the log-likelihood surface for the X-29 aircraft. The lower panels show the difference between the log-likelihoods for the two aircraft as a function of yaw and pitch. The lower right panel shows this difference as an image, such that regions where the log-likelihood for the incorrect aircraft type is higher than that for the correct aircraft type are shown in gray, and regions where the log-likelihood for the correct aircraft type is higher than that for the incorrect aircraft type are shown in white. The true orientation is marked by an "X" in the lower right panel.

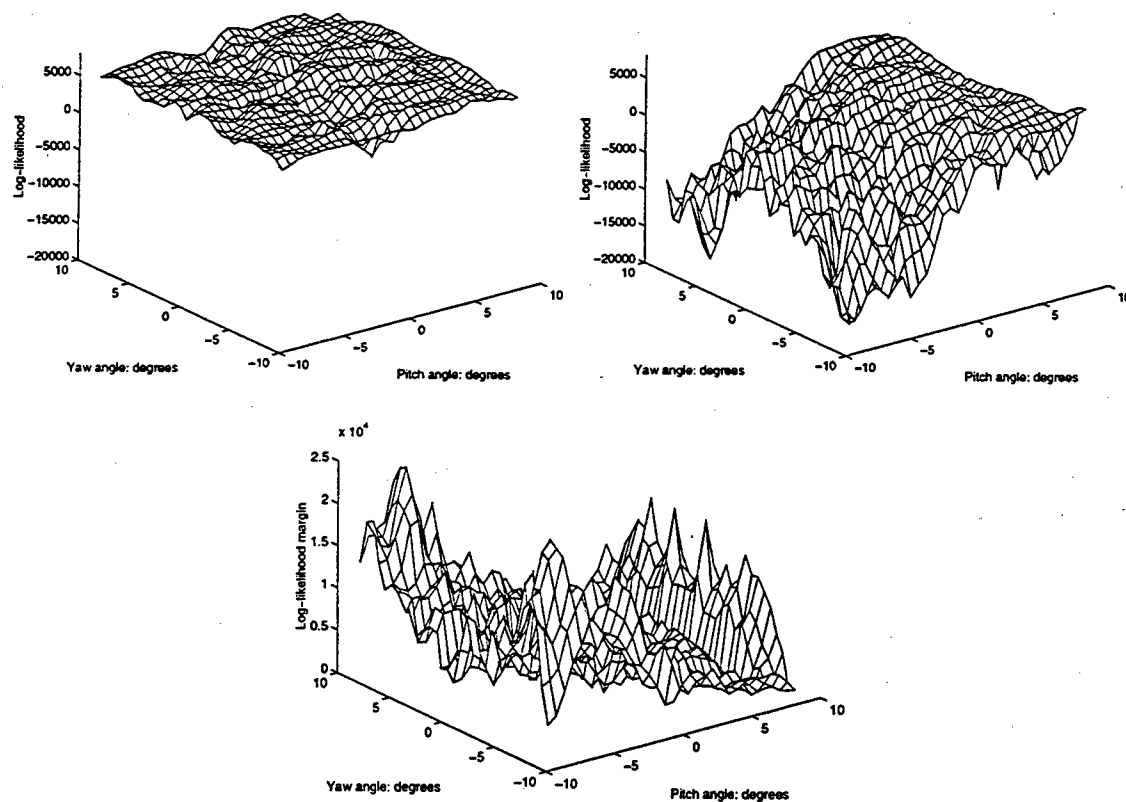


Figure 3.19: Conditionally Gaussian model log-likelihood functions for a two-aircraft recognition test. The observation was the range-profile for the F-117 aircraft at -4.34 degrees yaw and 4.37 degrees pitch. The upper left panel shows the log-likelihood surface for the F-117 aircraft at varying orientations. The upper right panel shows the log-likelihood surface for the X-29 aircraft. The lower panel shows the difference between the log-likelihoods for the two aircraft as a function of yaw and pitch.

lead to incorrect recognition. Closer inspection of the data plotted in the two previous figures reveals that, under either model, the log-likelihood for the correct target in the vicinity of the true orientation is larger than that for the incorrect target by a margin of roughly 4000.

The tests conducted so far have shown consistently superior performance for the conditionally Gaussian model over the deterministic model, but these tests have considered data from only a single polarization channel, namely vertical polarization during both transmission and reception. In the next subsection, we consider the use of data from multiple polarizations, describe variations on the conditionally Gaussian model to incorporate polarization diversity, and perform further tests to investigate if the recognition performance is improved.

3.3 Incorporation of Multiple Polarizations

In the development thus far, we have considered a radar signal to be a single, complex-valued function of time. Of course an actual radar signal is in fact an electromagnetic wave, and a complete characterization of the signal should be in terms of the propagation of the electromagnetic field. In this section, we discuss the concept of wave polarization, and extend the conditionally Gaussian model to polarimetric radar data. We proceed to repeat the recognition tests presented previously for the resulting polarimetric model, using two different forms for the polarimetric covariance.

3.3.1 Wave Polarization

For a plane wave traveling in free space, the electric field and magnetic field vectors remain orthogonal to each other and to the direction of propagation, while their lengths and directions vary in time and space. The polarization state of the wave is defined by the locus of points traced out by the electric field vector. Let a point be chosen at a fixed distance along the radar line-of-sight, and consider the variation of the electric field vector in the plane perpendicular to the line-of-sight at this point. For a fully polarized wave, the tip of the electric field vector traces out an ellipse in this plane over time [40]. If time is fixed, then the tip of the electric field vector forms an elliptical spiral in space. This general case is known as elliptical polarization. The polarization ellipse is defined by the length of the elliptical axes and by the tilt

angle defining the angular orientation of the elliptical axes with respect to the spatial coordinate axes.

Circular polarization is the case where the axes of the polarization ellipse are of equal length. Left-circular and right-circular polarization correspond to the cases where the rotation of the electric field vector with time is clockwise and counter-clockwise, respectively [40].

Linear polarization is the case where one of the axes of the polarization ellipse is of zero length, so that the electric field vector varies along a line segment. If in addition the tilt angle is zero, then the polarization is either horizontal or vertical.

Radar reflections from a target will in general depend upon the polarization state of the incident wave. An important consequence of the physics of plane wave polarization is that a wave with arbitrary elliptical polarization may be represented as the sum of two appropriately chosen linearly polarized waves. Thus, the response of a target to plane waves of arbitrary polarization can be summarized by the horizontally and vertically polarized reflections in response to horizontally and vertically polarized incident waves [46]. This characterization is typically defined in terms of a scattering matrix. The scattering matrix formulation for monochromatic illumination at frequency f can be defined

$$\begin{bmatrix} E_h^r(f) \\ E_v^r(f) \end{bmatrix} = \begin{bmatrix} S_{hh}(f) & S_{hv}(f) \\ S_{vh}(f) & S_{vv}(f) \end{bmatrix} \begin{bmatrix} E_h^i(f) \\ E_v^i(f) \end{bmatrix}, \quad (3.46)$$

where $[E_h^i(f) \ E_v^i(f)]^T$ are the horizontally and vertically polarized incident fields, and $[E_h^r(f) \ E_v^r(f)]^T$ are the reflected fields. For the entries in the scattering matrix, the first subscript denotes the polarization state during reception and the second subscript denotes the polarization during transmission. Note that the effects of range to the target and antenna gain have been ignored in this formulation.

Equation (3.46) suggests a simple technique to measure the scattering matrix for a target: transmit a horizontally polarized wave, measure the horizontally and vertically polarized responses, and repeat with a vertically polarized transmission. The resulting matrix provides a complete characterization of the scattering properties of the object in question at a particular orientation [16]. The inverse Fourier transform of each of the entries in the scattering matrix yields a time-domain representation for the scattering properties of the object, and can be collectively thought of as a set of four reflectivity profiles $b(\tau)$ for the object for each of the linear polarization states.

The entries in the scattering matrix share a deterministic interdependence defined by the physics of the measurement process. For monostatic radar, in which the transmitter and receiver are co-located, the cross-polarization terms are equal, $S_{hv}(f) = S_{vh}(f)$ [46, 42, 13]. Additionally, if the scattering matrix for an object at a particular orientation is known, then the scattering matrix resulting from the rotation of the object about the line-of-sight through an arbitrary angle can be found. To demonstrate this, imagine that the receive antennae consist of a pair of long wires that are horizontally and vertically aligned. Then rotation of the object about the line-of-sight through angle θ is equivalent to rotation of the antennae through $-\theta$. The scattering matrix resulting from this rotation is given by [46, 42]

$$\begin{bmatrix} S_{hh}(\theta) & S_{hv}(\theta) \\ S_{vh}(\theta) & S_{vv}(\theta) \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} S_{hh} & S_{hv} \\ S_{vh} & S_{vv} \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \quad (3.47)$$

where the explicit dependence on frequency has been suppressed here. Note that, if circular polarization is used during transmission and reception, then the resulting scattering matrix is completely invariant to the rotation in question. This can be seen by noting that for any such rotation, all points on the target retain their range coordinate, and the restriction to circular polarization guarantees that the strength of the field impinging on any point on the target is unchanged. The relationship demonstrated here is used in Chapter 4 to derive a transformation from an arbitrary object orientation defined by three Euler angles to a pair of angles defining a class of range-profile-equivalent orientations.

The use of horizontal and vertical polarization during transmission and reception makes it possible to measure range-profiles for a target over four polarization channels; let these four signals be denoted as $\{r_{hh}(t), r_{hv}(t), r_{vh}(t), r_{vv}(t)\}$. Due to the equivalence in the scattering matrix shown above, $r_{hv}(t) = r_{vh}(t)$ so that a complete characterization of the polarization response of the target is possible by measuring $\{r_{hh}(t), r_{hv}(t), r_{vv}(t)\}$. The HRR simulator XPATCH provides precisely this type of data for the target. In the next subsection, we consider application of the conditionally Gaussian model to the case of polarimetric observations, and repeat the recognition tests to investigate if there is an improvement in performance as a result of using polarimetric data.

3.3.2 Extension of Conditionally Gaussian Model

We now consider extension of the conditionally Gaussian model to the case of polarimetric data. The return in range bin l is now given by a 3-vector of returns for each of the polarization states,

$$\mathbf{r}_p[l] = \begin{bmatrix} r_{hh}[l] \\ r_{hv}[l] \\ r_{vv}[l] \end{bmatrix}. \quad (3.48)$$

These vectors are stacked to form the new observation vector \mathbf{r}_p . Conditioned on knowledge of the target type a and orientation θ , this vector is modeled as complex Gaussian with mean $\mu_p(\theta, a)$ and covariance matrix $\mathbf{K}_p(\theta, a)$. We continue to model the returns from different range bins as statistically uncorrelated. Thus the structure of \mathbf{K}_p is seen to be block-diagonal, where the 3×3 diagonal blocks each represent the covariance for the measurements in a single range bin,

$$\mathbf{K}_p[l] = E\{(\mathbf{r}_p[l] - \mu_p[l])(\mathbf{r}_p[l] - \mu_p[l])^\dagger\}. \quad (3.49)$$

The log-likelihood computation in equation (3.44) is unchanged by the generalization to polarimetric data, except for this change in the covariance matrix. The implementation of this log-likelihood expression in our recognition tests will depend on the structure assumed for $\mathbf{K}_p[l]$.

Diagonal Polarimetric Covariance

A simple structure that may be assumed for $\mathbf{K}_p[l]$ is diagonal, implying that the returns for different polarization states in the same range bin are uncorrelated. This is a physically inaccurate assumption, as any reflecting surface on the target having significant extent in both the horizontal and vertical directions will produce both horizontally and vertically polarized responses. However, it provides a simple way to extend the recognition tests performed earlier. The results of this extension are shown in Figure 3.20. Comparison of Figures 3.7 and 3.20 indicates that the inclusion of data from multiple polarization states resulted in an increase in the log-likelihood margin for the correct target near the true azimuth angle by approximately a factor of 2. For three of the panels, corresponding to the observation of the fire truck, the school bus, and the m1 tank model, the inclusion of multiple polarizations had no significant effect on the performance as a function of azimuth angle. Comparison of




Figure 3.20: Conditionally Gaussian model log-likelihood functions for four-vehicle recognition tests using polarimetric data and diagonal polarimetric covariance. In the upper left panel, the range-profile for the fire truck at 19.95 degrees azimuth was selected as an observation, and the conditionally Gaussian log-likelihood was computed for each vehicle for azimuth angles from 10 to 30 degrees. The lower left panel shows the result when the observation is the school bus at 25 degrees azimuth. The two right panels show the results when the observations are each of the two tank models at 15 degrees azimuth. In each panel, the solid line shows the log-likelihood for the fire truck, the dotted line represents the m1 tank, the dashed line represents the school bus, and the dot-dash line represents the t1 tank.

the lower right panels, however, indicates that for observation of the t1 tank, the range in azimuth over which correct recognition is achieved was decreased by the inclusion of multiple polarizations, due to the deep trough in the log-likelihood function for the t1 tank in the vicinity of 18 degrees azimuth. Apparently, this change in performance was due to variability in the hh and hv range-profiles for the t1 tank for azimuth angles near 18 degrees that is not present in the vv range-profiles.

Non-diagonal Polarimetric Covariance

To further investigate the use of polarimetric data, we generalize the model by allowing the hh and vv returns from the same range bin to be correlated. Thus, we now consider the covariance matrix $\mathbf{K}_p[l]$ to have the following structure

$$\mathbf{K}_p[l] = \begin{bmatrix} \sigma_{hh}[l] & 0 & \rho[l] \\ 0 & \sigma_{hv}[l] & 0 \\ \rho^*[l] & 0 & \sigma_{vv}[l] \end{bmatrix} \quad (3.50)$$

where $\sigma_{hh}[l]$, $\sigma_{hv}[l]$, and $\sigma_{vv}[l]$ are the auto-covariance terms for the individual polarization states used in the diagonal model above, and

$$\rho[l] = E\{(r_{hh}[l] - \mu_{hh}[l])(r_{vv}[l] - \mu_{vv}[l])^*\} \quad (3.51)$$

is the cross-covariance between the hh and vv returns. A covariance matrix of this form is proposed in [13] for polarimetric modeling of radar returns from clutter sources such as vegetation and snow-covered terrain. The model is formulated by applying a second order Born approximation to the wave equations describing the scattering behavior for a medium with randomly fluctuating permittivity. The first order terms in the Born approximation correspond to excitation by the incident wave of a particle within the medium that directly produces a backscattered radiation. The second order terms correspond to excitation by the incident wave of a particle within the medium that produces radiation exciting other particles in the medium that in turn produce backscattered radiation. This covariance structure is also used in [59, 57, 58] to investigate speckle reduction and target detection performance in polarimetric SAR imagery.

We implemented the conditionally Gaussian log-likelihood function (3.44) using the covariance structure (3.50), and repeated the recognition tests shown previously under this model. The results of these tests are shown in Figure 3.21. Comparison

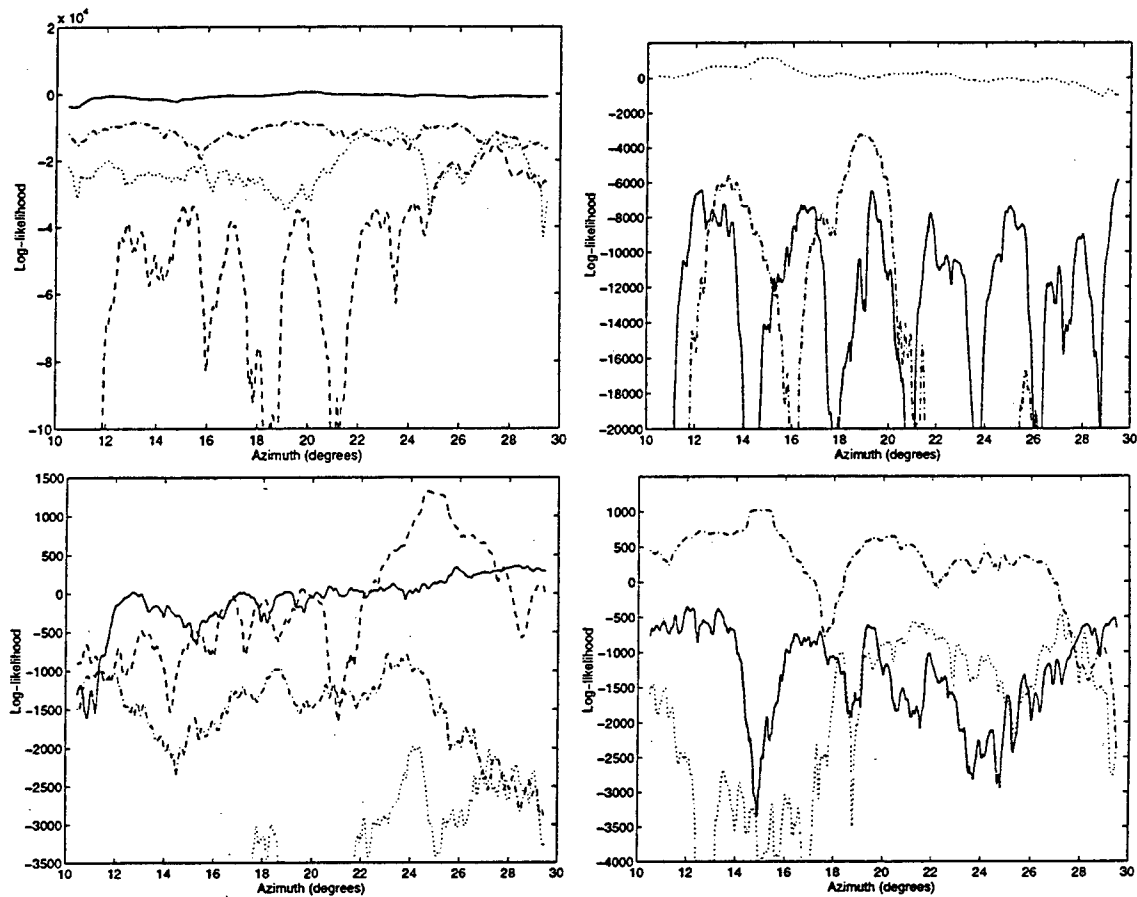


Figure 3.21: Conditionally Gaussian model log-likelihood functions for four-vehicle recognition tests using polarimetric data and non-diagonal polarimetric covariance. In the upper left panel, the range-profile for the fire truck at 19.95 degrees azimuth was selected as an observation, and the conditionally Gaussian log-likelihood was computed for each vehicle for azimuth angles from 10 to 30 degrees. The lower left panel shows the result when the observation is the school bus at 25 degrees azimuth. The two right panels show the results when the observations are each of the two tank models at 15 degrees azimuth. In each panel, the solid line shows the log-likelihood for the fire truck, the dotted line represents the m1 tank, the dashed line represents the school bus, and the dot-dash line represents the t1 tank.

of Figures 3.20 and 3.21 indicates that the use of the non-diagonal polarimetric covariance resulted in an improvement in performance over the diagonal polarimetric covariance for each of the four cases. In general, the change had little effect on the log-likelihood for the correct target, but it resulted in a decrease in the log-likelihood of the incorrect targets. This is most clearly seen in the lower right panels. In this case, the log-likelihood for the correct target, the t1 tank, is increased only slightly, whereas the log-likelihoods for the fire truck and the m1 tank are dramatically decreased for azimuth angles near the truth. As a result, correct recognition is achievable over a much wider range in azimuth when using the non-diagonal polarimetric covariance as compared to the diagonal polarimetric covariance.

3.4 Stochastic Reflectivity Model for Correlated Range-Profiles

The models for high resolution radar data that were examined in Sections 3.2 and 3.3 assumed that successive range-profiles collected from a given target are statistically independent. This assumption may not be valid for real HRR data. Consider a moving target that is illuminated by a sequence of radar pulses in rapid succession. For aircraft or ground vehicles, the target kinematics must be consistent with an appropriate physical model such as Newton's laws governing motion, so that the sequence of orientations assumed by the target at the illumination times will be correlated with one another. If the change in the orientation of the target is small enough between illuminations, then consecutive range-profiles of the target will likewise be correlated. The deterministic and conditionally Gaussian models do not allow for the incorporation of correlated range-profiles into the data likelihood. In this section, we give a brief outline of a model for range-profiles that attempts to account for this correlation.

There are several possible ways in which a stochastic model for the reflectivity may be defined. In this section, we describe a model based on a glint plus a diffuse component. The glint component incorporates the deterministic reflectivity profile as modeled by signature prediction tools such as XPATCH, while the diffuse component models the reflectivity as a random process as in the WSSUS model presented in Section 3.1.3

The reflectivity density of the target at location ρ is given by

$$c(\rho) = c_g(\rho)e^{j\psi(\rho)} + c_s(\rho), \quad (3.52)$$

where c_g is a glint component and c_s is a diffuse or speckle component. The phase ψ is a random variable with a known distribution that may depend on the angle between the line of sight and the surface normal at position ρ . The glint component is deterministic while the diffuse component is assumed to be a complex Gaussian random process that is spatially white. That is,

$$E\{c_s(\rho)\} = 0 \quad (3.53)$$

$$E\{c_s(\rho)c_s^*(\rho')\} = \Sigma(\rho)\delta(\rho - \rho'), \quad (3.54)$$

where $\Sigma(\rho)$ is the scattering function. This model is an extension of the standard WSSUS model described previously. It follows closely the model described by Shapiro for laser radar illuminations [69, 68, 70], and is a spatially-varying version of the Rician model [86].

Note that if $\psi = 0$ and $c_s = 0$, then this model reduces to the deterministic model of Section 3.2.1. If $c_g = 0$ and $c_s \neq 0$, then the reflections are purely diffuse, and the received signal is a zero-mean Gaussian random-process.

Let the reflectivity density-function on the surface, $c(\rho)$, be mapped to a reflectivity density-function in three dimensions via multiplication by an appropriate Dirac impulse function. Denote the resulting density function by $c(\mathbf{x})$ in a body-centered reference frame, where $\mathbf{x} = [x \ y \ z]^T$. The received signal can then be represented by

$$s(t; \theta) = \sqrt{E_t} \int S_T(f) e^{j2\pi ft} \int e^{j\frac{4\pi f}{c} \mathbf{u}^T \mathbf{x}} c(\mathbf{x}) d\mathbf{x} df, \quad (3.55)$$

where $S_T(f)$ is the Fourier transform of the transmitted signal, and

$$\mathbf{u} = \begin{bmatrix} \sin \phi_1 \cos \phi_2 \\ \sin \phi_1 \sin \phi_2 \\ \cos \phi_1 \end{bmatrix} \quad (3.56)$$

is the body-centered line-of-sight unit vector. Note that we have suppressed explicit dependence on the target type a in this formulation. Note also that equation (3.56) is of the same form as equation (3.16) for the range-spread target model, except that here the integration of the reflectivity density is performed over the entire target surface, and $c(\mathbf{x})$ is allowed to have non-zero mean.

The covariance function for $r(t)$ depends on the covariance function for $c(\mathbf{x})$. In order to get specific results, we assume that $\psi = 0$. Then, the covariance function

depends only on the covariance function for $c_s(\mathbf{x})$ and hence only on the scattering function. In this case, the range-profile $r(t)$ is a complex Gaussian random process with mean

$$\mu(t; \theta) = \sqrt{E_t} \int S_T(f) e^{j2\pi ft} \int e^{j\frac{4\pi f}{c} \mathbf{u}^T \mathbf{x}} c_g(\mathbf{x}) d\mathbf{x} df. \quad (3.57)$$

The cross-covariance function for a pair of range-profiles $r_{m_1}(t)$ and $r_{m_2}(t)$ is given by

$$\begin{aligned} K_r(t, t'; \theta_{m_1}, \theta_{m_2}) &= E_t \int \int S_T(f) S_T^*(f') e^{j2\pi(f t - f' t')} \\ &\quad \times \int e^{j\frac{4\pi}{c}(f \mathbf{u}^T - f' \mathbf{u}'^T) \mathbf{x}} \Sigma(\mathbf{x}) d\mathbf{x} df df' \\ &\quad + N_0 \delta_{m_1 m_2} \delta(t - t'), \end{aligned} \quad (3.58)$$

where \mathbf{u} and \mathbf{u}' are the unit vectors for the two illuminations.

The log-likelihood function for a sequence of range-profiles $\{r_k(t)\}$ may be derived from (3.58) using the results from Chapter 2 of [87]. This gives a log-likelihood conditioned on the target type and on the sequence of target orientations: $L(R_k | \Theta_k, a)$. Note that in the deterministic and conditionally Gaussian models, the log-likelihood of a sequence of range-profiles equals the sum of the individual log-likelihoods, whereas in this stochastic model it is crucial to include the terms quantifying the correlation of the sequence. For a single range-profile, the log-likelihood determined by this model, $L(r | \theta, a)$, can be used directly in place of the deterministic log-likelihood derived in the previous section, for cases in which the glint plus speckle model is appropriate. To date, we have not implemented this model for the purpose of comparison of recognition performance with the other models. However, we anticipate that this model may give better results when applied to real data than the other models, if the correlation between range-profiles at distinct orientations is significant.

3.5 Conclusions

This chapter has discussed a number of basic concepts in high resolution radar modeling, and has continued by introducing two candidate models for the range-profile, and testing the potential for recognition under each model using simulated HRR data. The various tests considered in this chapter demonstrated successful recognition using range-profiles simulated by XPATCH under both the deterministic and the conditionally Gaussian models. The performance seen under the conditionally Gaussian model

was generally superior to that under the deterministic model, in that the conditionally Gaussian model provided a log-likelihood margin for the correct target over all other targets over a wide range of azimuth angles. To extrapolate that recognition from real HRR data is possible using these models would require establishing a number of other facts:

- that the behavior seen in this example is reproduceable for other targets at other orientations,
- that recognition is possible in the presence of noise,
- that recognition is possible when attempting to distinguish between many more target types, some of which are macroscopically very similar, and
- that real HRR data from a given target are accurately modeled by the output of XPATCH.

The first two of these items have been investigated on a limited basis, as reported in Chapter 5, with a generally positive result. The third item could be investigated given datasets for enough targets, but has not been studied in this thesis. The last item is indeed the most critical of all, but lies outside the scope of this thesis. The developers of XPATCH are continuing to refine their product and the target models to be used with it, in an ongoing effort to establish the validity of the range-profiles produced.

Chapter 4

Orientation Estimation Using HRR Data

As discussed in Chapter 3, a significant difficulty encountered in recognition from range-profile data is the extreme variability observed in range-profiles for small changes in target orientation. Our method of dealing with this variability is to jointly estimate the target type and orientation sequence. In this chapter, we consider several issues that arise when attempting to estimate target orientation from range-profiles. We begin with an introduction to the representation of orientation in \mathbb{R}^3 using a 3×3 rotation matrix that takes its value in the special orthogonal group $\text{SO}(3)$. We proceed to describe global errors in orientation estimates derived from HRR data that can arise from target symmetry and invariance of the data to rotation about the radar line-of-sight. A method of resolving the rotational ambiguity using linearly polarized range-profiles is discussed. Performance bounds on orientation estimation are described in terms of a Hilbert-Schmidt bound on the estimation error, following the development of these bounds provided in [76]. These bounds are evaluated using simulated data to compare the performance under the two models for HRR data described in Chapter 3 and for two different HRR sensors. We conclude the chapter with a discussion of the work by Cutaia [23] and Srivastava [76] incorporating kinematic priors on aircraft motion into the ATR problem.

4.1 The Rotation Matrix

For an arbitrary object in \mathbb{R}^3 , the orientation of the object is defined by the relationship between its body axes and some specified set of coordinate axes. For example, in

describing the orientation of aircraft it is customary to define orientation in terms the rotation necessary to align a set of inertial axes with the body axes of the aircraft [27]. Formally, a rotation in \mathbb{R}^3 is defined as a linear transformation that preserves the inner product. Let $\mathbf{x}_B \in \mathbb{R}^3$ and $\mathbf{x}_I \in \mathbb{R}^3$ be the coordinates of a point on the object in the body-centered and inertial reference frames, respectively. Then orientation is defined by the matrix \mathbf{R} such that

$$\mathbf{x}_B = \mathbf{R}\mathbf{x}_I. \quad (4.1)$$

The requirement that rotation preserve inner product constrains \mathbf{R} to be an orthogonal matrix:

$$\mathbf{x}_B^T \mathbf{x}_B = \mathbf{x}_I^T \mathbf{x}_I \iff \mathbf{R}^T \mathbf{R} = \mathbf{I}. \quad (4.2)$$

Clearly, $|\mathbf{R}| = \pm 1$. In order that \mathbf{R} represent the rotation of a solid object in \mathbb{R}^3 , the determinant must be positive

$$|\mathbf{R}| = 1. \quad (4.3)$$

An orthogonal matrix with negative determinant corresponds to the composition of rotation and reflection about an axis. For example, if the transformation

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

were applied to a left-handed person, it would leave him right handed. So the orientation of an object in \mathbb{R}^3 can be described by a 3×3 rotation matrix, which is constrained to be an orthogonal matrix with determinant 1, also known as a special orthogonal matrix. The set of 3×3 rotation matrices has a group structure [76]; this group is known as the special orthogonal group 3, and is denoted by $\text{SO}(3)$.

It is also the case that the group $\text{SO}(3)$ can be identified with the direct product of the unit sphere and the unit circle, $\mathbf{S}^2 \times \mathbf{S}^1$ [76]. To see this, let the rotation matrix \mathbf{R} have rows $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3$. Then the matrix equation (4.2) can be viewed as nine scalar equations constraining $\{\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3\}$. From the diagonal we have,

$$\mathbf{r}_1 \mathbf{r}_1^T = \mathbf{r}_2 \mathbf{r}_2^T = \mathbf{r}_3 \mathbf{r}_3^T = 1 \quad (4.4)$$

Thus, each row vector is constrained to the unit sphere. The off-diagonal terms of (4.2) require

$$\mathbf{r}_1 \mathbf{r}_2^T = \mathbf{r}_1 \mathbf{r}_3^T = \mathbf{r}_2 \mathbf{r}_3^T = 0 \quad (4.5)$$

Suppose \mathbf{r}_1 is chosen to be an arbitrary unit vector. Then \mathbf{r}_2 is constrained to the great circle orthogonal to \mathbf{r}_1 . Once \mathbf{r}_2 is chosen from this circle, there are only two unit vectors \mathbf{r}_3 orthogonal to both \mathbf{r}_1 and \mathbf{r}_2 ; one is the negative of the other. Of these two vectors, only one will satisfy $|\mathbf{R}| = 1$.

The parameterization of rotations as elements of $\mathbf{SO}(3)$ is extremely useful from the point of view of modeling rigid target motion, in that [76]

- it provides a simple rule of composition through matrix multiplication,
- it allows for the computation of an average of orientations as the normalized average of the rotation matrices, and
- it allows for a definition of distance between two orientations in terms of the Hilbert-Schmidt distance between the rotation matrices.

However, there are other parameterizations. For example, an arbitrary rotation in \mathbb{R}^3 can be achieved by rotation about an appropriately chosen axis through an appropriately chosen angle. Additionally, a rotation can be parameterized by three Euler angles representing the composition of three rotations about the body axes. In the next section, we make use of the Euler representation to consider a class of orientations that produce equivalent range-profiles in the case of circular polarization.

4.2 Orientation Invariance of HRR data

Suppose we are given the problem of estimating the orientation of a known object from a single observation of the range profile. Errors in the estimates result from a variety of sources. Previous work by Sworder, et al. [82] identifies both local and global errors arising out of the problem of target identification through matching a noisy observation with a stored replica, when the data to be compared are 2-D silhouettes of the target. For range profile data, the problem is analogous. Local and global errors may arise due to the variability of the range profile with the orientation of the object and the effects of noise in the observations. Later in this chapter we will consider fundamental limits on estimation accuracy as a result of these errors. Other types of global errors result from deterministic sources and correspond to the fact

that different orientations of the target yield equal range profiles under our model. We use the terms symmetry and ambiguity to describe these deterministic sources of error. In this section we discuss these errors and justify the dependence of the range-profile on only two Euler angles.

For this discussion, we will define the geometry of our experiment consistent with the treatment of rigid body dynamics used in the tracking literature (see Friedland [27]). The target can be described by a set of points in one of two spaces: an inertial frame of reference which is assumed fixed, and a body-centered frame of reference that is independent of orientation. Each frame of reference is described by three orthogonal axes with respect to a common origin at the center of the target. The inertial x -axis is assumed to be the radar line-of-sight; the body-centered x , y , and z axes point to the front, right side, and bottom of the target respectively.

The orientation of a target is defined by the orientation of its body axes with respect to the inertial axes. An arbitrary orientation can be described by a set of three Euler angles (ϕ, θ, ψ) , called roll, pitch, and yaw angles. Previously in this thesis we have used θ to denote the complete orientation of the object, but here it represents only the pitch angle. For an aircraft target, changing the yaw angle corresponds to rotating the target about the z_B -axis, or rotation in the wing plane. Changing the pitch corresponds to rotation about the y_B -axis, or tipping the nose of the aircraft up or down. Changing the roll corresponds to rotating about the x_B -axis, or rolling the wings around the nose axis. The Euler angles define an ordered triple of rotations from the inertial to the body-centered frame of reference as follows:

$$\begin{bmatrix} x_B \\ y_B \\ z_B \end{bmatrix} = T_{IB} \begin{bmatrix} x_I \\ y_I \\ z_I \end{bmatrix} \quad (4.6)$$

$$T_{IB} = R_x(\phi)R_y(\theta)R_z(\psi) \quad (4.7)$$

where

$$R_x(\phi) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\phi) & \sin(\phi) \\ 0 & -\sin(\phi) & \cos(\phi) \end{bmatrix} \quad (4.8)$$

$$R_y(\theta) = \begin{bmatrix} \cos(\theta) & 0 & -\sin(\theta) \\ 0 & 1 & 0 \\ \sin(\theta) & 0 & \cos(\theta) \end{bmatrix} \quad (4.9)$$

$$R_z(\psi) = \begin{bmatrix} \cos(\psi) & \sin(\psi) & 0 \\ -\sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (4.10)$$

The ordering of the rotations is implied by the above matrix product: yaw then pitch then roll. This ordering is physically desirable because it guarantees that each rotation does not affect the previous rotations. For example, consider performing a pitch followed by a yaw. After the pitch, the wing plane is no longer aligned with the x_I - y_I plane. Therefore, the yaw rotation will change the angle between the x_B axis and the x_I - y_I plane, thereby effectively changing the pitch of the aircraft with respect to the inertial coordinate system.

Many targets of interest such as aircraft exhibit symmetry with respect to the body-centered x - z plane. Assume that the reflectivity of the target is given by the summed contribution of a finite number of scattering sites, and the range profile results from projection of the individual reflectivities along the line of sight and sorting into range bins. Under these conditions the range profile will be invariant to reflection in the inertial x - z plane yielding the potential global symmetry error. In other words, a symmetric target at a yaw angle of ψ will have the same distribution of reflectivity along the line-of-sight as the same target at a yaw angle of $-\psi$. In one of the ATR simulations described in Chapter 5, an example is shown where the estimation algorithm produces a global error resulting from target symmetry. Such errors are unresolvable without including the prior arising from the state space description of the target kinematics.

Even in the absence of target symmetry, there is a potential for global ambiguity. This is because, if circular polarization is used during transmission and reception, the range profile is invariant to rotation about the inertial x -axis. The justification for this invariance can be seen by noting that, under rotation about the x_I -axis, all points on the target retain their range coordinate, any self-shadowing of target structures is unchanged, and the restriction to circular polarization guarantees that the strength of the impinging wavefront at any point on the target is also unchanged. As discussed in Chapter 3, if linear polarization is used, then the range profile is not invariant to line-of-sight rotation, but the set of range-profiles obtained under any

two such rotations are deterministically related through (3.47). Global ambiguities are resolvable using either kinematic information or multiple linear polarizations.

Define an equivalence class to be the set of orientations of a target that produce identical range profiles under circular polarization. Given an orientation for a target, other orientations in the same equivalence class arise by rotating about the x_I axis. Since x_I is not in general aligned with any of the body axes, a three step process is employed.

1. Perform a series of rotations about major axes so that the x_I axis is aligned with one of the body axes. There are many possible solutions, but we will choose to simply invert the original rotations T_{IB} .

$$\begin{bmatrix} x_1 \\ y_1 \\ z_1 \end{bmatrix} = T_{IB}^{-1} \begin{bmatrix} x_B \\ y_B \\ z_B \end{bmatrix} = \begin{bmatrix} x_I \\ y_I \\ z_I \end{bmatrix} \quad (4.11)$$

2. Perform an arbitrary rotation about the x_I axis. Let the angle of rotation be α . Since the x_I and x_B axes are now aligned, the necessary rotation is simply $R_x(\alpha)$.

$$\begin{bmatrix} x_2 \\ y_2 \\ z_2 \end{bmatrix} = R_x(\alpha) \begin{bmatrix} x_1 \\ y_1 \\ z_1 \end{bmatrix} = R_x(\alpha) \begin{bmatrix} x_I \\ y_I \\ z_I \end{bmatrix} \quad (4.12)$$

3. Invert the rotations applied in step 1.

$$\begin{bmatrix} x_3 \\ y_3 \\ z_3 \end{bmatrix} = T_{IB} \begin{bmatrix} x_2 \\ y_2 \\ z_2 \end{bmatrix} = T_{IB} R_x(\alpha) \begin{bmatrix} x_I \\ y_I \\ z_I \end{bmatrix} \quad (4.13)$$

The ordering of rotations presented above makes good physical sense. However, it is possible to achieve all orientations regardless of this ordering. For example, suppose we chose to perform roll first.

$$\begin{bmatrix} x_B \\ y_B \\ z_B \end{bmatrix} = R_y(\theta) R_z(\psi) R_x(\phi) \begin{bmatrix} x_I \\ y_I \\ z_I \end{bmatrix} \quad (4.14)$$

This choice of ordering makes it easy to determine all orientations in an equivalence class. If we now apply our three step process:

$$\begin{bmatrix} x_3 \\ y_3 \\ z_3 \end{bmatrix} = R_y(\theta)R_z(\psi)R_x(\phi)R_x(\alpha) \begin{bmatrix} x_I \\ y_I \\ z_I \end{bmatrix} \quad (4.15)$$

Since $R_x(\phi)R_x(\alpha) = R_x(\phi + \alpha)$, this equation has the interesting implication that if we set $\alpha = -\phi$, the associated equivalence class for each orientation has at least one member with $\phi = 0$. Thus, every equivalence class is uniquely identified by a pair of angles θ' and ψ' which are the pitch and yaw of that member of the equivalence class with zero roll. The rotations, in standard order, in equation (4.13) define an equivalence class. By our previous argument there must be some angle α and some pair (θ', ψ') such that the following equation is satisfied:

$$R_y(\theta')R_z(\psi')R_x(-\alpha) = R_x(\phi)R_y(\theta)R_z(\psi) \quad (4.16)$$

The three angles that solve this matrix equation must satisfy each of the nine scalar equations contained within. However, if we are interested only in the angles (θ', ψ') that identify the equivalence class for a given orientation, we may consider the three scalar equations from the first column, which are independent of α .

$$\cos(\theta')\cos(\psi') = \cos(\theta)\cos(\psi) \quad (4.17)$$

$$-\sin(\psi') = -\cos(\phi)\sin(\psi) + \sin(\phi)\sin(\theta)\cos(\psi) \quad (4.18)$$

$$\sin(\theta')\cos(\psi') = \sin(\phi)\sin(\psi) + \cos(\phi)\sin(\theta)\cos(\psi) \quad (4.19)$$

The unique solution to these equations defines a mapping from a complete orientation to a pair of angles that identify an equivalence class:

$$\begin{aligned} \psi' &= \arcsin(\cos(\phi)\sin(\psi) - \sin(\phi)\sin(\theta)\cos(\psi)) \\ \theta' &= \arctan\left(\frac{\sin(\phi)\sin(\psi) + \cos(\phi)\sin(\theta)\cos(\psi)}{\cos(\theta)\cos(\psi)}\right) \end{aligned} \quad (4.20)$$

The set of equivalence classes defined by the two angles (θ', ψ') can be identified with the unit sphere S^2 . The points of closure on this sphere correspond to the sets of orientations for which $\psi' = \pm \frac{\pi}{2}$. These two values for ψ' each define an equivalence class, regardless of the angle θ' .

The analysis presented in this section has defined a class of orientations that produce equivalent range-profiles when circular polarization is used during transmission and reception. This ambiguity may be resolved if tracking information is available, or if circular polarization is not used. In the next section we investigate how the deterministic dependence identified in equation (3.47) can be exploited to resolve the global ambiguity.

4.3 Estimation of Line-of-Sight Rotation Using Polarimetric Data

The range-profiles from a target under circular polarization are invariant to rotation of the target about the radar line-of-sight. This invariance does not extend to the case of linear polarization, but the range-profiles resulting from an arbitrary rotation about the line-of-sight are deterministically related through equation (3.47). This has the implication that, even for the case of linear polarization, libraries of range profiles need only be stored over the set of angles (θ', ψ') that define equivalence classes. The range-profiles for any orientation within the equivalence class can be generated on-line through (3.47). In this section, we consider taking advantage of this deterministic relationship to estimate the rotation about the line-of-sight, thereby eliminating the global ambiguity. We thank Paritosh Tyagi for his contribution to the following derivation.

Let \mathbf{X} represent a single-frequency measurement of the scattering matrix for the target,

$$\mathbf{X} = \begin{bmatrix} x_{hh} & x_{hv} \\ x_{vh} & x_{vv} \end{bmatrix}. \quad (4.21)$$

It is assumed that \mathbf{X} represents the measurement when the target assumes a neutral orientation with respect to the line-of-sight. For consistency with the last section, let the orientation be given by yaw and pitch angles (θ', ψ') with $\alpha = 0$. Let \mathbf{Y} represent

the measurement obtained for some other rotation α . Then we have,

$$\begin{bmatrix} y_{hh} & y_{hv} \\ y_{vh} & y_{vv} \end{bmatrix} = \begin{bmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{bmatrix} \begin{bmatrix} x_{hh} & x_{hv} \\ x_{vh} & x_{vv} \end{bmatrix} \begin{bmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{bmatrix}. \quad (4.22)$$

As discussed in Chapter 3, the cross-polarization terms are equal in \mathbf{X} and \mathbf{Y} . Let the independent elements of \mathbf{X} and \mathbf{Y} be arranged into column vectors as

$$\mathbf{x} = \begin{bmatrix} x_{hh} \\ x_{hv} \\ x_{vv} \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} y_{hh} \\ y_{hv} \\ y_{vv} \end{bmatrix}. \quad (4.23)$$

If we carry out the matrix multiplication on the right side of equation (4.22), we find that \mathbf{y} is related to \mathbf{x} through

$$\mathbf{y} = \mathbf{A}(\alpha) \mathbf{x}, \quad (4.24)$$

where the matrix \mathbf{A} is given by

$$\mathbf{A}(\alpha) = \begin{bmatrix} \cos^2 \alpha & \sin 2\alpha & \sin^2 \alpha \\ -\frac{1}{2} \sin 2\alpha & \cos 2\alpha & \frac{1}{2} \sin 2\alpha \\ \sin^2 \alpha & -\sin 2\alpha & \cos^2 \alpha \end{bmatrix}. \quad (4.25)$$

This matrix has determinant 1, but is not orthogonal. However, the inverse has a very similar form,

$$\mathbf{A}^{-1}(\alpha) = \begin{bmatrix} \cos^2 \alpha & -\sin 2\alpha & \sin^2 \alpha \\ \frac{1}{2} \sin 2\alpha & \cos 2\alpha & -\frac{1}{2} \sin 2\alpha \\ \sin^2 \alpha & \sin 2\alpha & \cos^2 \alpha \end{bmatrix}. \quad (4.26)$$

Additionally, the matrix $\mathbf{A}(\alpha)$ and its inverse may be decomposed as

$$\mathbf{A}(\alpha) = \mathbf{A}_1 \cos 2\alpha + \mathbf{A}_2 \sin 2\alpha + \mathbf{A}_3 \quad (4.27)$$

$$\mathbf{A}^{-1}(\alpha) = \mathbf{A}_1 \cos 2\alpha - \mathbf{A}_2 \sin 2\alpha + \mathbf{A}_3 \quad (4.28)$$

where

$$\mathbf{A}_1 = \begin{bmatrix} \frac{1}{2} & 0 & -\frac{1}{2} \\ 0 & 1 & 0 \\ -\frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix}$$

$$\begin{aligned} \mathbf{A}_2 &= \begin{bmatrix} 0 & 1 & 0 \\ -\frac{1}{2} & 0 & \frac{1}{2} \\ 0 & -1 & 0 \end{bmatrix} \\ \mathbf{A}_3 &= \begin{bmatrix} \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix}. \end{aligned} \quad (4.29)$$

Let the data \mathbf{x} be modeled as a 0-mean complex Gaussian random vector with covariance matrix \mathbf{K} . Then \mathbf{y} is also complex Gaussian with mean 0 and covariance

$$\mathbf{K}_y = \mathbf{A}(\alpha) \mathbf{K} \mathbf{A}^T(\alpha). \quad (4.30)$$

The log-likelihood for \mathbf{y} given α is

$$L(\mathbf{y}|\alpha) = -\ln |\mathbf{K}_y| - \mathbf{y}^\dagger \mathbf{K}_y^{-1} \mathbf{y} \quad (4.31)$$

$$= -\ln |\mathbf{K}| - \mathbf{y}^\dagger \mathbf{A}^{-T}(\alpha) \mathbf{K}^{-1} \mathbf{A}^{-1}(\alpha) \mathbf{y} \quad (4.32)$$

where † denotes complex conjugate transposition. A necessary condition for the maximum likelihood estimate α_{ml} is that it must satisfy

$$\frac{\partial L(\mathbf{y}|\alpha)}{\partial \alpha} = 0. \quad (4.33)$$

Performing the differentiation, we arrive at

$$B_1 \sin 4\alpha + B_2 \cos 4\alpha + B_3 \sin 2\alpha + B_4 \cos 2\alpha = 0, \quad (4.34)$$

where the coefficients are given by

$$B_1 = 2\mathbf{y}^\dagger (\mathbf{A}_1^T \mathbf{K}^{-1} \mathbf{A}_1 - \mathbf{A}_2^T \mathbf{K}^{-1} \mathbf{A}_2) \mathbf{y} \quad (4.35)$$

$$B_2 = 2\mathbf{y}^\dagger (\mathbf{A}_2^T \mathbf{K}^{-1} \mathbf{A}_1 - \mathbf{A}_1^T \mathbf{K}^{-1} \mathbf{A}_2) \mathbf{y} \quad (4.36)$$

$$B_3 = \mathbf{y}^\dagger (\mathbf{A}_3^T \mathbf{K}^{-1} \mathbf{A}_1 - \mathbf{A}_1^T \mathbf{K}^{-1} \mathbf{A}_3) \mathbf{y} \quad (4.37)$$

$$B_4 = \mathbf{y}^\dagger (\mathbf{A}_3^T \mathbf{K}^{-1} \mathbf{A}_2 - \mathbf{A}_2^T \mathbf{K}^{-1} \mathbf{A}_3) \mathbf{y}. \quad (4.38)$$

Equation (4.34) may be reduced to the form

$$C \sin(4\alpha + \xi) = \sin 2\alpha \quad (4.39)$$




Figure 4.1: Plots of $\sin 2\alpha$ and $\frac{1}{2} \sin \left(4\alpha + \frac{\pi}{6} \right)$.

where

$$C = -\sqrt{\frac{B_1^2 + B_2^2}{B_3^2 + B_4^2}} \quad (4.40)$$

$$\xi = \tan^{-1} \left(\frac{B_2}{B_1} \right) - \tan^{-1} \left(\frac{B_4}{B_3} \right). \quad (4.41)$$

The left and right sides of equation (4.39) are plotted in Figure 4.1 for $C = \frac{1}{2}$ and $\xi = \frac{\pi}{6}$. In general there will be 4 solutions for $|C| < 1$ and 8 solutions for $|C| > 1$, only one of which can be α_{ml} . Equation (4.34) can also be transformed into an equivalent polynomial equation of order four through the substitution

$$u = \sin 2\alpha \quad \sqrt{1 - u^2} = \cos 2\alpha, \quad (4.42)$$

but this approach does not lead to an analytical solution in the general case. For a particular instance of \mathbf{y} and \mathbf{K} , this technique yields four solutions that may be individually tested for maximizing $L(\mathbf{y}|\alpha)$. Alternatively, numerical approaches such as Newton's method may be employed to determine the solutions.

4.4 Performance Metrics for Orientation Estimation

The problem of automatic target recognition has received wide attention in recent years, and a variety of solutions have been proposed. Many of the techniques using range-profile data are described in Chapter 2. Given this diversity of approaches, it is desirable to have a basis upon which various candidate algorithms can be compared. In Chapter 3, we examined simulated range-profiles for various targets, and observed that the data exhibited extreme variability for small changes in target aspect angle. It is in part this observation that has motivated posing the ATR problem in terms of joint estimation of target type and orientation. Clearly, our ability to recognize targets from range-profiles will be impacted by the accuracy to which the target orientation can be estimated from the available data. In this section, we consider performance bounds on orientation estimation using HRR data. These bounds provide specific model selection criteria for this problem, and will allow us to evaluate different types of HRR sensors for their utility in estimating orientation.

The orientation of a target takes its value on a curved space. For a parameterization of the orientation in terms of a rotation matrix this space is $\text{SO}(2)$ or $\text{SO}(3)$. The question arises as to how one can evaluate estimators on these spaces; the standard Cramer-Rao analysis assumes that the parameter lies in a patch in \mathbb{R}^n . Following closely chapter 9 of Srivastava [76], we consider a Hilbert-Schmidt bound on the expected error in orientation estimation. This provides a lower bound on the performance of any estimator using a given set of data. We will restrict our attention to the case of estimators on $\text{SO}(2)$, but the generalization to $\text{SO}(n)$ is straightforward and is given in [76].

To begin the development, we consider a 2×2 rotation matrix $\mathbf{O}(\theta) \in \text{SO}(2)$ defining the orientation of a target relative to a high resolution radar. For example, let \mathbf{O} define the orientation of a ground target illuminated by an airborne radar platform. For a fixed elevation angle, θ is the azimuth of the target relative to the radar. The Hilbert-Schmidt norm of $\mathbf{O}(\theta)$ is defined as

$$\|\mathbf{O}(\theta)\|_{HS} = \sqrt{\sum_{i,j=1}^2 O_{ij}^2}. \quad (4.43)$$

Because \mathbf{O} is a rotation,

$$\mathbf{O}(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}, \quad (4.44)$$

and $\|\mathbf{O}(\theta)\|_{HS}^2 = 2$. Furthermore, for any two rotations \mathbf{O}_1 and \mathbf{O}_2 ,

$$\|\mathbf{O}_1 - \mathbf{O}_2\|_{HS}^2 = 4 - 4(\cos \theta_1 \cos \theta_2 + \sin \theta_1 \sin \theta_2) \quad (4.45)$$

$$= 4 - 4 \cos(\theta_1 - \theta_2) \quad (4.46)$$

$$= 4 - 2 \operatorname{tr}(\mathbf{O}_1 \mathbf{O}_2^T). \quad (4.47)$$

Thus, the Hilbert-Schmidt norm provides a natural distance metric between elements of $\mathbf{SO}(2)$.

Given a true orientation for the target \mathbf{O}_{true} and any estimate $\hat{\mathbf{O}}$, the squared Hilbert Schmidt norm $\|\mathbf{O}_{true} - \hat{\mathbf{O}}\|_{HS}^2$ gives a measure of the squared error in the estimate. Let \mathbf{r} denote an observed range-profile from which the estimate $\hat{\mathbf{O}}$ is computed, and let $p(\mathbf{r}|\mathbf{O})$ denote the likelihood function for \mathbf{r} given \mathbf{O} . Also, let $p(\mathbf{O})$ be the prior on the orientation. A Hilbert-Schmidt estimator is defined as the estimator that, given \mathbf{r} , minimizes the conditional mean under the posterior of the squared estimation error as defined by the Hilbert-Schmidt norm [76],

$$\hat{\mathbf{O}}_{HS} = \operatorname{argmin}_{\mathbf{O} \in \mathbf{SO}(2)} \int_{\mathbf{SO}(2)} \|\mathbf{O} - \mathbf{O}'\|_{HS}^2 p(\mathbf{O}'|\mathbf{r}) \gamma(d\mathbf{O}'), \quad (4.48)$$

where γ is the base measure on $\mathbf{SO}(2)$. As indicated by Srivastava, this estimator is a minimum mean-squared error estimator with the mean defined using the Hilbert-Schmidt norm, and it is optimal in the sense that for any estimator, the expected squared error is lower bounded by the squared error achieved by the Hilbert-Schmidt estimator [76]. To relate this estimator to conventional notions of mean-squared error, consider the form for the Hilbert-Schmidt norm-squared given by

$$\|\mathbf{O} - \hat{\mathbf{O}}\|_{HS}^2 = 4 - 4 \cos \theta_d \quad (4.49)$$

where θ_d is the difference between the true azimuth and the estimated azimuth. Expanding the cosine in a Taylor's series and retaining terms of order two and lower, we have

$$\|\mathbf{O}_1 - \mathbf{O}_2\|_{HS}^2 \approx 2\theta_d^2. \quad (4.50)$$

Thus, for sufficiently small errors θ_d , the squared error in the Hilbert-Schmidt sense is twice the squared error in the azimuth estimate $\hat{\theta}$.

The Hilbert-Schmidt bound on the expected squared error in estimation of a rotation in $\text{SO}(2)$ is defined as [76]

$$HSB = E \left\{ \int_{\text{SO}(2)} \|\hat{\mathbf{O}}_{HS} - \mathbf{O}\|_{HS}^2 p(\mathbf{O}|\mathbf{r}) \gamma(d\mathbf{O}) \right\}, \quad (4.51)$$

where γ is the base measure on $\text{SO}(2)$, and the outer expectation is over the ensemble of observations \mathbf{r} . We will compute this bound as a function of the signal-to-noise ratio in the observations for different models for range-profiles. This provides a criterion for model selection for the problem of orientation estimation. We will also compare the bounds for different examples of HRR sensors, allowing us to evaluate the optimal performance achievable using data from each sensor.

The computation of the Hilbert-Schmidt bound involves evaluating the two expectations given in (4.51). Following the example in [76], we will assume that the underlying true orientation is fixed, and generate a simulated range-profile for the target at this orientation. In this case, the outer expectation is approximated by averaging over multiple realizations of the noise added to form the observations. The inner expectation over candidate orientations \mathbf{O} will be evaluated by numerical integration. Additionally, we assume a uniform prior $p(\mathbf{O})$. With these simplifications, the algorithm to compute the Hilbert-Schmidt bound is given by Srivastava [76]. The matrix \mathbf{A} is defined as the conditional mean of \mathbf{O} under the likelihood function $p(\mathbf{r}|\mathbf{O})$ and is approximated by

$$\mathbf{A} = \frac{\sum_{i=0}^{N-1} \mathbf{O}(\theta_i) p(\mathbf{r}|\mathbf{O}(\theta_i))}{\sum_{i=0}^{N-1} p(\mathbf{r}|\mathbf{O}(\theta_i))}, \quad (4.52)$$

where the angles θ_i are equally spaced around the circle, $\theta_i = \frac{2\pi i}{N}$. The Hilbert-Schmidt estimate is found by normalizing \mathbf{A} to unit determinant, which projects this matrix into $\text{SO}(2)$,

$$\hat{\mathbf{O}}_{HS} = \frac{\mathbf{A}}{\sqrt{|\mathbf{A}|}}. \quad (4.53)$$

The squared error in this estimate is computed using the Hilbert-Schmidt norm-squared, and the HSB is found by averaging the squared error over multiple realizations of the noise process.

The Hilbert-Schmidt bound was evaluated using simulated range-profile data under each of the HRR models discussed in Chapter 3. We thank Matthew Cooper




Figure 4.2: Hilbert-Schmidt bound on expected squared error in azimuth estimation for the “M1” tank model viewed at 30 degrees azimuth. The solid line represents the performance under the deterministic model and the dashed line represents the performance under the conditionally Gaussian model.

for his extensive work in developing the software that was used in computing the bounds. For this evaluation, we have used the L-band data from the URISD for the “M1” tank model. The range-profile for this vehicle at 30.6 degrees azimuth was selected as the observation. Libraries developed for each of the models were provided to the algorithm for the purpose of evaluating the likelihood function. The library for the deterministic model consisted of URISD range profiles for the tank at 360 discrete azimuth angles equally spaced around the circle. The library for the conditionally Gaussian model consisted of 240 means and variances computed over 3 degree patches in azimuth. The patches were equally spaced around the circle, overlapping neighboring patches by 1.5 degrees. The HSB algorithm described above was implemented using these libraries. Two hundred independent realizations of the additive noise were generated for noise variances between 6.25×10^{-6} and 64 square meters. The corresponding values for the signal-to-noise ratio vary from -18 dB to 50 dB.

The resulting performance bounds are shown in Figure 4.2. In this figure, a value of 0 for the HSB indicates perfect azimuth estimation in that the estimate is equal to the truth. A value of 4 for the HSB indicates completely unreliable azimuth estimation, in the sense that this is the performance that would be achieved on average for an estimator whose estimates are random with uniform density on the circle. Note that under both models, nearly perfect performance is achieved for high signal-to-noise ratios, and the performance degrades rapidly with increasing noise.

The figure indicates equivalent performance under either model over the entire range of signal-to-noise ratios considered. Based on the results of Chapter 3, we might expect that the performance under the deterministic model would be better. The reasoning here is that we observed that the likelihood function under the deterministic model exhibited higher curvature near the true orientation than under the Gaussian model. Thus, a Cramer-Rao type analysis would indicate that more accurate orientation estimation is possible using the deterministic model. However, we have only used 360 discrete orientations in evaluating the bound for the deterministic model, and have therefore significantly undersampled the likelihood functions shown in Chapter 3.

A second evaluation of the Hilbert-Schmidt bound was conducted to compare the performance using two different HRR sensors. For one of the sensors, we used the L-band data from the URISD for the "M1" tank model. These data correspond to a microwave radar with center frequency 1.5 GHz and bandwidth 739.5 MHz, resulting in a range resolution of 20.3 centimeters. For the other sensor, we used the X-band data from the URISD for the "M1" tank model. These data correspond to a microwave radar with center frequency 10 GHz and bandwidth 1.4775 GHz, resulting in a range resolution of 10.15 centimeters.

In this case, both libraries consisted of 360 simulated range-profiles for target azimuths equally spaced around the circle. The deterministic model likelihood function was used exclusively. The range-profile for a target azimuth of 30 degrees was selected from each data set as the observation, and the HSB computation was performed using 200 independent realizations of the noise process for variances between 6.25×10^{-6} and 64 square meters. The corresponding values for the signal-to-noise ratio vary from -18 dB to 50 dB. The resulting performance bounds are shown in Figure 4.3. This plot does not show the results for signal-to-noise ratios higher than 30 dB, as the bound for both sensors is zero at these noise levels. The figure indicates essentially equivalent performance using either sensor, except at very high noise levels where the L-band sensor shows better performance. This example indicates that the performance is not improved by either the higher bandwidth or the higher center frequency associated with the X-band sensor. We might expect both of these factors to affect performance: the bandwidth determines the range resolution and the center frequency determines the rate at which range-profiles decorrelate for small changes in orientation. However, the results seen here were compiled for a single target at a fixed true orientation, and the performance may be different if these parameters are varied.




Figure 4.3: Hilbert-Schmidt bound on azimuth estimation error for the “M1” tank model viewed at 30 degrees azimuth. The solid line represents the performance using simulated data corresponding to an L-band radar, and the dashed line represents the performance using simulated data corresponding to an X-band radar.

Additionally, we are computing the Hilbert-Schmidt bound using only 360 discrete orientations in each case, and it was shown in Chapter 3 that many more orientation samples would be required to accurately reproduce the likelihood function.

The results presented in this section have made use of a fundamental metric for estimators taking their value on $SO(2)$ or $SO(3)$, the Hilbert-Schmidt bound, to compare limits on the performance of orientation estimation using two different models for the same sensor, and two different sensors under the same model. In the first case, the bound indicated equivalent azimuth estimation performance under either model. In Chapter 5 results are presented indicating that a conditionally Gaussian model is more appropriate for ATR. In the second case, the bound did not distinguish either the L-band or the X-band radar as superior in terms of orientation estimation performance. Clearly what is needed in both cases is a more extensive analysis that would consider varying true orientations, and would bound the performance on joint estimation of orientation and target type. Such analysis comprises a major goal for future research on this problem.

4.5 Kinematic Priors for ATR

As discussed in Chapter 1, a primary factor making the problem of recognition of targets from range-profile data difficult is the extreme variability in the observed data for small changes in target aspect angle. It has therefore consistently been

our view that observation of a single range-profile is insufficient for recognition, and this has led us to pose the problem in terms of jointly estimating the sequence of target orientations and the target type from a sequence of observed range-profiles. Given this framework for the problem, successful recognition is crucially dependent on a kinematic prior on the orientation sequence that is derived from an appropriate dynamical model for the target motion. Such a prior is also a crucial element of the system proposed in [31, 47, 76] for the general multi-target ATR problem. As demonstrated by Cutaia [23], such a prior can be used to estimate target class directly from tracking data, where target class refers to a broad characterization such as “jet fighter” or “large cargo aircraft.” Additionally, a kinematic prior can provide considerable information about the likely position and orientation for a target at the time of a future observation. This is the approach used by Srivastava [76], and is directly applicable to ATR using HRR data. In this section, we discuss the work of these authors addressing the incorporation of prior information on target motion into algorithms for ATR, and go on to discuss the need for kinematic priors to properly address the problem of recognition from HRR data.

4.5.1 Recognition Using Kinematic Measurements

Cutaia [23] explores the problem of recognition of aircraft targets by exploiting prior information on aircraft-class-dependent kinematics. A kinematic model for aircraft motion is derived from Newton’s second law; this consists of twelve nonlinear differential equations relating the positions and orientations of the target and their derivatives to the forces and torques acting on the target. The model includes aircraft-class-specific physical parameters that can be used to discriminate targets. These parameters such as wing area, weight, and moments of inertia vary by several orders of magnitude depending on the aircraft class in question. The model is linearized about an equilibrium condition and decomposed into two systems of differential equations that separately describe the longitudinal and lateral target motion, respectively. In the linearized model, the inputs are pilot commands representing deflections of the control surfaces. These commands take their value from a discrete set that is also aircraft-class-dependent and are modeled as a Markov chain. The resulting state-variable model for aircraft motion has the form of a discrete-time jump-linear system.

Given this model, the tracking problem becomes a conditional-mean state estimation problem, and the recognition problem becomes a problem of choosing one

of N systems that maximizes the posterior probability density conditioned on the observations. The observations consist of noisy measurements of the aircraft position and orientation in an inertial coordinate system, obtained from a tracking radar and an imaging sensor respectively. The optimal solution to either the tracking or recognition problem is shown to be intractable, as the posterior probability density conditioned on the measurements consists of a Gaussian mixture where the number of member densities grows exponentially with the time index k . Suboptimal solutions from the literature including the generalized pseudo-Bayesian (GPB) and interacting multiple-model (IMM) algorithms are investigated. Bounds are derived on the performance degradation resulting from these suboptimal solutions in terms of the increase in the probability of incorrect recognition, the L_1 distance between the optimal and suboptimal posterior density functions, and relative hybrid entropy between the optimal and suboptimal posteriors.

Cutaia presents an aircraft recognition example in which a jet fighter aircraft performing a 4g turn is to be distinguished from a cargo aircraft performing a 1g turn. The data used for recognition are noisy measurements of the position and orientation for one of the targets. A jump-linear discrete-time system model is developed for each aircraft; this consists of specifying the set of allowable pilot input states and the associated input-state transition probability matrix for each aircraft. One of the two aircraft is selected as the true target, and observations are generated for this target performing a maximally capable turn. The IMM algorithm is used to track the target under hypothesis A, that the true target is the jet fighter, and hypothesis B, that the true target is the cargo aircraft. The simulation results clearly indicate that acceptable tracking performance is achieved only when the correct target model is used, and that recognition is possible using the ratio of the conditional log-likelihoods under each model. This work has demonstrated that, using only tracking information and aircraft-class-dependent kinematic models, an aircraft can be recognized to the level of the class, if not to the level of an individual target type. That is not to say, however, that this work does not have a significant role to play in the general ATR problem. If the majority of possible target types can be eliminated from consideration based on kinematic measurements alone, then incorporation of these priors into ATR systems could significantly improve performance.

4.5.2 Kinematic Priors in Multisensor ATR

Srivastava [76] extends this analysis by explicitly incorporating geometry into the ATR problem. Each target in the scene is represented by a CAD model for the target surface referred to as a static generator. The position and orientation of a particular target in the scene at a fixed time are represented by the translation and rotation applied to the associated generator. For ground targets the translation p lies in \mathbb{R}^2 and the rotation O lies in $\text{SO}(2)$, whereas for airborne targets the translation lies in \mathbb{R}^3 and the rotation in $\text{SO}(3)$. A dynamic generator is defined as the combination of the parameters for a particular target in the scene over a time interval: the position and orientation at the start of the interval, the static generator representing the target type which remains fixed over the interval, and the position and orientation at the end of the interval. Rules for concatenation of dynamic generators are established such that the sequence of positions and orientations produced are meaningful, in that the position and orientation at the end of each time interval are equal to the position and orientation at the start of the next interval, and the target type remains fixed. The resulting concatenated sequence forms a linear graph and represents a target track, i.e. samples of the target trajectory at discrete times. The complete configuration space is built in a hierarchical manner from individual tracks by combining tracks over different target types, different track lengths, and finally differing numbers of tracks in the scene. A distance metric on the complete configuration space is established so that the notion of minimum mean square error estimation on this space is meaningful.

A kinematic model for aircraft targets derived from Newtonian mechanics is also used here. This model is distinct from that used by Cutaia in two important ways: the equations are not linearized about an equilibrium point, and the input forces and torques are not modeled in terms of specific pilot commands to the control surfaces of the aircraft. The model is discretized to yield a set of nonlinear difference equations relating the positions, orientations, forces and torques at discrete times j , and the forces and torques are modeled as independent Gaussian random variables. The resulting discrete-time equations of motion are analyzed to induce a prior on the future position and orientation of a target, given previous values.

This prior is combined with the data likelihood at the time of the next observation to provide the posterior. By deriving priors explicitly in $\mathbb{R}^3 \times \text{SO}(3)$, the natural parameter space for the target kinematic state, the range of likely positions and orientations for the target is considerably reduced, thereby reducing the computational burden required to produce estimates from the observations. These are precisely the

priors needed to couple into the work comprising this thesis. Srivastava presents two ATR simulations using the kinematic prior. The first simulation addresses the problem of joint detection, tracking, and recognition of a single maneuvering target using data from both a passive, narrowband tracking array and an optical imaging sensor. The second simulation demonstrates joint detection and tracking of two maneuvering targets using the same tracking array and two of the imaging sensors, one dedicated to each of the targets.

4.6 Kinematic Models for ATR Using HRR

The approach to ATR described in this thesis is significantly different from other approaches proposed in the literature. A very significant shortcoming of most conventional approaches to the ATR problem in a radar context is that they do not consider the role that kinematic state estimation has to play in recognition. These researchers assume that the orientation is known or resort to exhaustive searches over some range of possible orientations to achieve recognition, but do not take advantage of the knowledge that the sequence of orientations assumed by the target must be consistent with a model for the target motion. In contrast, we maintain that estimation of target dynamics is required for recognition from range-profiles. The resulting prior on target motion may not be precisely that used by Cutaia or Srivastava, but it should be closely related in that it is derived from a Newtonian kinematic model. An important area of focus during the early part of the research program resulting in this thesis was a collaborative effort with Srivastava to incorporate range-profile data into ATR simulations using the kinematic prior described in [76]. One of the resulting simulations is presented in Chapter 5. A comprehensive simulation demonstrating joint tracking and recognition from HRR data and using kinematic prior on the orientations has not been completed, but this is an important area for future research.

Both Cutaia and Srivastava focus primarily on the case in which the sensors are fixed on the ground and the targets are aircraft. The Newtonian model from which they derive their priors is appropriate in this context. In the simulations presented in Chapter 5, we will consider an air-to-air scenario and an air-to-ground scenario. Defining kinematic priors appropriate to these contexts would require extensions or simplifications of those described above. In the case of an air-to-air encounter, kinematic models for both the sensor platform and the target could be developed, thereby

producing priors on the positions and orientations of both aircraft relative to an inertial coordinate system. This problem can be greatly simplified by assuming that the sensor platform is equipped with on-board sensors capable of providing accurate measurements of position and orientation; a Global Positioning System receiver and a gyroscope, for example. In the air-to-ground context, there is the additional problem of mapping the kinematic models for aircraft to describe the motion of wheeled vehicles. The resulting model may be a simplification of the aircraft model to planar motion in \mathbb{R}^2 , as suggested by Srivastava. One important feature of the motion of ground vehicles that is not achievable by most fixed-wing aircraft is the ability to simply stop moving and start again at unannounced times. Appropriate models for ground target motion must account for this behavior.

In the simulations presented in Chapter 5, we assume that no information about target position is available to the estimation algorithm. Thus, the kinematic models used by Srivastava and Cutaia are not applicable. We have instead chosen to apply a very simple linear model to each of the Euler angles describing the target orientation,

$$\dot{\theta}_j(t) = q_j(t), \quad (4.54)$$

where the orientation rates q_j are assumed to be slowly varying with respect to the rate at which range-profiles are collected. It must be noted that we in no way consider this model to be the "right" model for this problem. The right models are those used by Srivastava and Cutaia, with the possible extensions to the air-to-air and air-to-ground scenarios described above. However, this simple model has allowed us to demonstrate in a simulation environment the potential for joint orientation tracking and recognition from range-profile data. The model is useful only for those situations in which the target trajectory through orientation space is sufficiently well-behaved and the data collection rate is sufficiently high that the variation of the Euler angles is approximately linear over the time required to collect a sequence of range-profiles.

To conclude, it is clear that priors on aircraft motion derived from kinematic models for the targets greatly enhances the potential for recognition. When using high resolution radar data for recognition of dynamic targets, observation of a single range-profile is insufficient for recognition. Estimation of a sequence of target orientations and target type from a sequence of range-profiles is crucially dependent on an appropriate kinematic prior of the variety used by Cutaia and Srivastava.

4.7 Conclusions

This Chapter has investigated various issues relating to the estimation of target orientation from high resolution radar data. Potential global errors in orientation estimates derived from range profiles were described. These errors result both from target symmetry and from invariance of the data to rotation about the line-of-sight. A potential technique for resolving the rotational ambiguity using linearly polarized measurements was discussed. Performance bounds on orientation estimation in the form of Hilbert-Schmidt bounds on the expected estimation error were described, and they were implemented to compare the performance under two models for HRR data and two different HRR sensors. Finally, a discussion of the importance of incorporating prior information on target motion was given, concluding that such priors can enhance recognition. In the next chapter, we describe simulations that have been constructed to demonstrate the potential for joint tracking and recognition using range-profiles.

Chapter 5

Simulations

In Chapter 1, we introduced a methodology for joint tracking and recognition of multiple targets using multiple sensors. Rather than developing a system to implement this methodology in its entirety, we wish to test the potential for orientation estimation and recognition using HRR data. To this end, we have devised three experiments in which ATR algorithms are tested in a simulated environment. In all cases, we have made assumptions to simplify the problem being addressed from the general multiple-target ATR problem presented in Chapter 1. These assumptions allow us to focus on specific issues relating to HRR data.

5.1 Joint Tracking of Position and Orientation for a Single Maneuvering Aircraft Target

In our first simulation, we consider the case where it is known *a priori* that a single target of known type inhabits the scene throughout the simulation. The goal is to demonstrate capabilities for joint Bayesian tracking of position and orientation from multiple data sets. We implement an algorithm for ATR based on jump-diffusion processes in this case. This algorithm has been presented previously for ATR with high-resolution images [48] and for multiple-target direction-of-arrival tracking [78]. The scenario addressed in this simulation differs from the scenario for the general ATR problem in that a single target is present, the target has been successfully detected, the target does not leave the scene during the simulation, and the target type is known. We thank Dr. Anuj Srivastava and Mohammad Faisal for their work in developing the simulation presented in this section.

There are two remote sensors which observe the dynamic target: a low resolution tracking array of passive, narrowband sensors sensitive to the position of the target but insensitive to the target shape and orientation, and a high resolution radar. At sampling times, the tracking array produces a complex vector modeled as the superposition of a signal contribution and the surrounding noise [49, 67]. In this model, the signal component is given by the complex product $s(t)d(p(t))$ where $s(t)$ is the complex amplitude of the signal wavefront, and $d(p(t))$ is called the direction or steering vector encoding the target position $p(t)$ depending on the geometry of the sensor arrangement [49, 48]. Assuming additive, white, Gaussian noise and deterministic signal amplitudes the log-likelihood of tracking data, $y(1), y(2), \dots$, is

$$L_2(p(1), \dots, p(N)) = -\frac{1}{\sigma^2} \sum_{i=1}^N \|y[i] - d(p[i])\|^2. \quad (5.1)$$

The joint log-likelihood is formed as the sum of this expression and the HRR data log-likelihood derived earlier, resulting in a natural fusion of information collected at different sensors. For this simulation, we have implemented our algorithm using the HRR log-likelihood function for the deterministic model in (3.40). In a simulation presented later in this chapter, we use the log-likelihoods derived from both the deterministic model and the conditionally Gaussian model, and compare the performance under each.

The rigid body dynamics governing target motion are analyzed to derive a prior probability density on tracking and recognition parameters. This prior, combined with the likelihood of data collected at sensing platforms, provides the posterior likelihood for Bayesian optimization. Following standard rigid body analysis [27], the translational velocities $\mathbf{v}(s) = [v_1(s) \ v_2(s) \ v_3(s)]^T$ and orientation rates $\mathbf{q}(s) = [q_1(s) \ q_2(s) \ q_3(s)]^T$ satisfy the following set of differential equations:

$$\begin{aligned} \dot{v}_1(s) - q_3(s)v_2(s) + q_2(s)v_3(s) &= f_1(s) \\ \dot{v}_2(s) + q_3(s)v_1(s) - q_1(s)v_3(s) &= f_2(s) \\ \dot{v}_3(s) - q_2(s)v_1(s) + q_1(s)v_2(s) &= f_3(s) \\ \\ I_1 \dot{q}_1(s) - (I_2 - I_3)q_2(s)q_3(s) &= \Gamma_1(s) \\ I_2 \dot{q}_2(s) - (I_3 - I_1)q_1(s)q_3(s) &= \Gamma_2(s) \\ I_3 \dot{q}_3(s) - (I_1 - I_2)q_2(s)q_1(s) &= \Gamma_3(s), \end{aligned} \quad (5.2)$$

where $\mathbf{f}(s) \in \mathbb{R}^3$ is the vector of forces driving the target and $\mathbf{\Gamma}(s)$ is the vector of applied torques. Assuming a statistical model on the forcing functions, the statistics on target positions were derived using the differential operator associated with these equations [48]. A jump-diffusion algorithm combining the kinematic prior obtained from (5.2) and the tracking data log-likelihood in (5.1) was derived for the case where a high resolution imaging sensor is used in conjunction with the cross-array tracking sensor [48]. In this simulation, we have adapted the algorithm to incorporate HRR data in place of the imaging data. Estimation of target orientation from HRR data proceeds according to a deterministic search through a library of simulated range-profiles as given below. A target of interest was selected (the X29 aircraft) and range-profiles were generated off-line for this target at 412 discrete aspect angles using XPATCH. These data are referred to as the library. The aspect angles were chosen so as to be approximately equally spaced around the sphere, with approximately 10 degrees separating nearest neighbors. Only the range profiles corresponding to the vv polarization state were retained for use in this simulation. Note that although the rotational invariance described in Section 4.2 does not strictly apply to the case of linear polarization, for the purposes of this simulation our library contains range-profiles for orientations parameterized by only two Euler angles, namely yaw and pitch. An orientation described by two Euler angles in this manner is hereafter referred to as the target aspect angle. As discussed below, we avoid the problem of estimating a complete target orientation described by three Euler angles from data compiled for only two such angles by assuming knowledge of the target roll angle.

At each time step when orientation estimation is to be performed, the true orientation is obtained from the flight path and converted to the associated aspect angle through (4.20). The range-profile corresponding to the library entry nearest the true aspect angle is selected as the observation, with no additive noise. The initial orientation estimate is set equal to the true orientation at the first time step, and equal to the previous estimate thereafter. This corresponds to a simplified dynamics model in which the orientation is constant between successive illuminations. Refinement of orientation estimates is performed by a deterministic walk on the likelihood grid. From the initial estimate, the algorithm examines the nearest neighbor library entries, and proceeds to the aspect angle of highest likelihood. This process is repeated for a fixed number of steps, or until a local maximum is reached. The aspect angle at which the algorithm terminates is considered the estimate. In order to convert this




Figure 5.1: Results of orientation estimation for simulation 1. In the left panel, the sequences of true yaw angles and yaw estimates are plotted. The right panel shows similar plots for the pitch angles.

aspect angle back into a complete orientation estimate, we assume knowledge of the true roll angle at each time step.

The algorithm has been implemented on a networked system of a DECmpp 128×128 massively parallel processor for array data simulation, computations, and efficient loading and processing of the range-profile library, and a Silicon Graphics workstation for simultaneous display of results. The original flight path was generated using a flight simulator package to provide the parameters for data simulation, the path sampled 1200 times at the sampling rate of 30/sec. Each sample is the true position and orientation state of the aircraft at the sampling instant. At each sampling time a new data set is added and the simulation performed over a new posterior on an enhanced parameter space.

The results of this simulation are shown in Figure 5.1. In the left panel of this figure, the sequence of true yaw angles is shown by a smooth curve, while the sequence of yaw estimates appears as a staircase approximation to this curve. The right panel shows similar plots for the pitch angles. Note that no results are given for tracking of the target position. The algorithm used for position tracking for this simulation is identical to that presented in [48], and the results obtained here were consistent with those results. The large discontinuities in the sequences of orientation estimates and the error in these estimates are a direct result of the extremely coarse orientation grid used in this simulation. This figure is also indicative of the general success of this simulation; as long as the initial orientation estimate at each time step was within 10 to 20 degrees of the truth, the algorithm generally converged to

the maximum-likelihood estimate. It should be noted that this success is critically dependent on the fact that the observed range-profile is always identical to one of the library entries. It is unlikely that the algorithm would succeed if the observations were range-profiles generated for orientations not in the grid.

This simulation demonstrates the potential of our methodology for joint tracking the position and orientation of a dynamic aircraft target observed by ground-based radar systems. However, due to the coarse orientation grid used here, errors in yaw and pitch estimates are often 5 degrees or more. Additionally, the algorithm for orientation estimation used in this simulation takes no advantage of the knowledge that the sequence of orientations assumed by the target should be consistent with Newtonian mechanics, as given by (5.2). The analysis of simulated HRR data presented in Chapter 3 has indicated that, when using the deterministic model, orientation estimates accurate to within 1 degree or less may be required for acceptable recognition performance. In our next simulation, we incorporate a much finer grid spacing in orientation, as well as a simple dynamical model that corresponds well with the dynamics of our simulated target, thereby allowing for much more accurate estimation of yaw and pitch angles and for successful recognition using HRR data.

5.2 Joint Orientation Estimation and Recognition for a Closing Aircraft Target

We have developed a computer simulation in which an unknown aircraft target is encountered by an airborne sensor platform at long range. It is assumed that detection and tracking of the position of this target are performed by an external system. A sequence of simulated range-profiles is observed as the target moves relative to the sensor. The goal is to estimate the target orientation during each simulated illumination and to determine the target type. The target type is constrained to be one of two possible candidates, namely either the F-117 or X-29 aircraft. The target and sensor follow straight-line flight paths at constant velocities and maintain roll angles of 0 degrees (level flight) during the encounter. These assumptions imply that the orientation of the target is completely specified by its yaw and pitch angles relative to the sensor. We further assume that the target is nominally nose-on to the sensor, such that the yaw and pitch angles lie between ± 10 degrees during the encounter. The reasons for choosing this particular scenario include:

- The focus here is on the simulated range-profiles and their utility in estimating orientation and identifying targets. Thus, other factors such as aircraft dynamics have been greatly simplified.
- The scenario is representative of a typical air-to-air encounter of an unknown target.
- The limited extent of possible aircraft orientations allows for development of a library of range-profiles computed on a fine grid in orientation, while keeping the total size of the library reasonable.

We have developed an estimation algorithm to operate under this scenario. We thank Mohammad Faisal for his assistance in developing the simulation presented here. Libraries of range-profiles for the F-117 and X-29 aircraft were simulated via XPATCH at a discrete set of orientations for yaw and pitch angles each varying between -10 and 10 degrees. The yaw angle was sampled uniformly at 0.2 degree intervals, and the pitch angle was sampled at a yaw-dependent interval chosen to keep the arc length on the unit sphere between neighboring pitches approximately 0.2 degrees.

Values were chosen for the initial position, initial orientation, and velocity of the target, and were used to compute a sequence of true orientation states for the target during the simulation. A true target type was chosen, and XPATCH was used to simulate the sequence of observed range-profiles corresponding to the sequence of true orientations. Thus, in contrast to our first simulation, the libraries and observations used here were computed independently.

In the previous simulation, estimates of target orientation were computed individually for each illumination, with the final estimate from the previous illumination providing an initial estimate for the current illumination. This technique can be problematic, as once the algorithm produces a significant error in the estimated orientation, all future estimates will be affected by this error. The dynamics-based prior presented previously addresses this problem by constraining estimates of target position and orientation over time to be consistent with Newtonian mechanics, but also requires that position and orientation be estimated jointly. In this simulation, we estimate orientation and target type only, and we have therefore chosen a very simple dynamical model to relate past and present orientation estimates; namely, that the yaw and pitch angles defining the aircraft orientation vary linearly over a sliding window in time. Thus, a sequence of orientations over the window are completely

determined by the yaw and pitch angles at the window endpoints. The window length W is specified in terms of an integer number of illuminations (e.g. $W = 32$).

The estimation algorithm proceeds as follows. After loading range-profile libraries for both aircraft and the first W observed range-profiles, initial estimates of target type and the first W orientations are generated. This is accomplished in several steps.

1. A global search through the range profile library for the F-117 is performed, and the likelihood of observation number 1 is computed for each candidate orientation. The ten orientations having the highest computed likelihood are retained. This process is repeated using observation number W , producing 100 pairs of candidate orientations at the window endpoints.
2. A global search through this list is conducted. For each entry, linear interpolation is applied to yaw and pitch estimates from the window endpoints to generate a sequence of orientation estimates over the window. Each of these estimates is replaced by the nearest orientation that is a grid point in our libraries. A sequence of range profiles corresponding to the sequence of grid points is extracted from the library. These data and the sequence of observations are used to compute the joint likelihood of the sequence of grid points corresponding to the current list entry, and the process is repeated for the entire list. The list entry yielding the highest joint likelihood is selected as the sequence of orientation estimates, given that the target type estimate is the F-117.
3. The entire process is repeated using the range-profile library for the X-29, producing a second estimate of the orientation sequence and target type. Of these two, the estimate yielding the higher joint likelihood is selected as the final estimate of target type and orientation sequence for the first W illuminations.

Future orientation estimates are generated through a local search.

1. The window endpoints are shifted by one illumination, and initial orientation estimates at the new endpoints are found by linear interpolation from the previous estimates. These estimates are replaced by their nearest grid points.
2. The set of nearest neighbors to each endpoint estimate is found. Our grid is nearly rectangular, so that in most cases, each endpoint has four neighbors, found by incrementing or decrementing the yaw or pitch.




Figure 5.2: Illustration of local search algorithm. The orientation estimates at the window endpoints are constrained to the grid values for which our libraries contain range-profiles. The current estimates are each represented by a filled circle and their neighbors are each represented by an "X". For each of the 25 pairs of candidates, the joint likelihood of the current sequence of observations is evaluated, and the pair of maximum likelihood is selected as the updated estimate. This process is repeated until a local maximum of the joint likelihood surface is reached.

3. Consider the current estimate and its neighbors as candidate orientations, so that we have five candidates at each endpoint or 25 possible pairs of candidates, each defining an orientation sequence over the window. These candidates are illustrated in Figure 5.2. The joint likelihood of each sequence is evaluated, using the current target type estimate, and the maximum-likelihood candidate pair is selected as the updated estimate. This process is repeated until the new estimate is equal to the current estimate. In this manner, the algorithm will proceed from the initial estimates produced by linear interpolation to the nearest local maximum of the joint likelihood surface.

Recognition is performed during the initialization process described above, in the sense that a global search is performed through the libraries for both target types. To allow for the possibility that this search may have produced the incorrect target type, the algorithm considers the hypothesis that the alternate target type is correct, and evaluates the joint likelihood of the current orientation estimates under this hypothesis. The local search is again implemented to find the nearest local maximum of the likelihood surface. The joint likelihood of the new orientation sequence and the alternate target type hypothesis is compared to that of the current estimates, and if the new likelihood is greater, then the target type estimate is changed to the alternate target type, and the current orientation estimates are replaced by the new




Figure 5.3: A frame of graphics output captured from simulation 2. In the right panel, the true sequence of yaw and pitch angles are shown in gray, with yaw and pitch estimates for the first 32 illuminations shown in black. The first 32 observed range profiles are shown in the upper left panel, and the sequence of library elements corresponding to the sequence of orientation estimates are shown in the lower left panel. The center panel displays a model of the true target type at the true orientation above a model of the estimated target type at the estimated orientation, for illumination number 16.

sequence. This process is invoked at regular intervals defined by an integer number of illuminations.

The software implementation includes graphics routines for visualization of the data and true and estimated parameters. A single frame of graphic output illustrating the state of the algorithm immediately following the initialization procedure described above is presented in Figure 5.3. In the right panel, the gray curves indicate the true sequence of yaw and pitch angles assumed by the target during the simulation, while the black line segments indicate the yaw and pitch estimates for the first 32 illuminations. A sequence of 32 observed range profiles and the library elements corresponding to the current orientation estimates are displayed in the left panel. The center panel contains a rendering of the true and estimated target types at the true and estimated orientations, for illumination number 16.

True orientation states for the simulation were generated as follows. The target begins the simulation at a range of 12.5 miles, at a yaw angle of -4 degrees relative to the line-of-sight of the radar, and at a pitch angle of 4 degrees. The aircraft flies




Figure 5.4: Results of orientation estimation for simulation 2. No initialization or recognition were performed in this case. In the left panel, the sequences of true yaw angles and yaw estimates are plotted. The right panel shows similar plots of the pitch angles.

along a straight line at 500 miles per hour for 10 seconds. The orientation angles are computed at 200 sample times over the duration of this flight.

In our first test of this simulation, we provided the algorithm with the true target type and the true target orientations for the first 32 observations (the initialization procedure and recognition step were not performed). The results of this test are presented in Figure 5.4. The left panel of this figure shows the sequence of true yaw angles and the resulting estimates. The estimates shown here are taken from the midpoints of the line segments defining the final estimates of the orientation sequences for each window position. The right panel shows similar plots for the pitch angles. As in simulation 1, the angle estimates appear as a staircase approximation to the smooth curve depicting the true angles. The discontinuities in the sequences of estimated angles and the resulting estimation error are small as compared to simulation 1 due to the fine orientation grid used here.

In a second test of our algorithm, the initialization procedure and recognition step have been included. Results from this test are shown in Figure 5.5. As before, the left panel shows the true yaw angles and estimates, and the right panel shows the true pitch angles and estimates. It is clear from the left panel that our initialization procedure produced a very significant error in this case. This error is interesting in that it results from a symmetry present in our libraries, which in turn results from a physical symmetry of the aircraft. The F-117 and X-29, like most aircraft, exhibit left-right symmetry, or symmetry with respect to the x - z plane in body-centered coordinates. Thus, any such aircraft rotated through a yaw angle




Figure 5.5: Results of joint orientation estimation and recognition for simulation 2. Both initialization and recognition were performed in this case. In the left panel, the sequences of true yaw angles and yaw estimates are plotted. The right panel shows similar plots of the pitch angles.

of ψ will have the same arrangement of reflectivity versus range as when rotated through a yaw angle of $-\psi$. XPATCH produces nearly identical range profiles at these two orientations, resulting in a symmetry in our range profile libraries, and thereby resulting in the orientation ambiguity encountered here. Note that this ambiguity is present throughout the simulation in that the estimated yaw angle is consistently close to the negative of the true yaw angle. Such ambiguities are resolvable if tracking information is available, which is not the case here. Note that we have conducted other tests of this algorithm in which a sign error in the yaw estimate did not occur.

Recognition results for this simulation were excellent. The algorithm consistently picked the correct aircraft type during the initialization procedure. Also, each time the algorithm considered the hypothesis that the alternate target type was correct, this hypothesis was rejected. Such a result was expected in this case, as the library range profiles for the two aircraft exhibited observable differences over the entire range of possible orientations, for example in the number and location of peaks.

This simulation has demonstrated the potential for successful joint orientation tracking and target recognition using simulated range-profile data, using the deterministic model log-likelihood (3.40) and a simple constant orientation rate dynamical model for the target motion. The simulation provided generally positive results for orientation tracking and very positive results for recognition in simplified ATR environment in which the target orientation is restricted to a 20×20 degree orientation window, and the target type is limited to one of two candidates. Further improvements on this simulation would include considering more candidate target types,

incorporating the conditionally Gaussian model for the range-profiles, and testing the performance of the algorithm when the observations include additive noise. All of these items are included in our third simulation, described in the next section.

5.3 Joint Orientation Tracking and Recognition of Ground Targets

The simulations presented in the previous section demonstrated the potential for jointly estimating target type and a sequence of target orientations from a sequence of simulated range-profiles. While demonstration of this potential is of interest, the simulations have not presented results quantifying the performance of the algorithm in terms of the rate of successful recognition or the accuracy of orientation estimates generated, nor has there been any consideration of the effect of additive noise on estimation performance. Also, in both simulations described earlier in this chapter, the deterministic model for range-profiles was used exclusively, largely because these simulations were developed prior to conducting the investigation of the conditionally Gaussian model presented in Chapter 3. The simulation presented in this section attempts to address each of these issues. Specifically, for the simplified ATR scenario described below, we investigate the performance of an algorithm for joint orientation tracking and recognition from simulated HRR data, consider the effect of noise on performance, and implement the algorithm under both the deterministic and the conditionally Gaussian models, to determine if the promising results obtained for the conditionally Gaussian model as presented in Chapter 3 are reproduced when using this model for ATR. We thank Zoltan Bekker for his assistance in developing these simulations.

For these simulations, we consider joint tracking and recognition of ground vehicles from an airborne sensor platform. The data available for estimation of target orientation and recognition are sequences of simulated range-profiles drawn from the University Research Initiative Synthetic Dataset, described in Section 3.2.1. The target is a rigid body that moves along flat ground in a straight line at constant velocity. The sensor platform flies in a straight line at constant velocity, and is assumed to maintain a constant elevation angle with respect to the target. A high resolution radar sensor mounted on the platform remains pointed at the target throughout the encounter. Given these assumptions, the target orientation θ is determined solely by

the azimuth angle of the target relative to the sensor, which takes its value on the circle. For the development of our estimation algorithm, we make the assumption as in the previous simulation that the azimuth angle varies linearly over a sliding window in time.

5.3.1 Observations

A particular instance of our simulation is defined by selecting one of the four vehicles {fire truck, school bus, m1 tank, t1 tank} as the true target, selecting a fixed elevation angle and frequency band for the simulation, and choosing values for the parameters governing the relative motion of the target and sensor. These parameters are used to compute the azimuth angle of the target relative to the sensor at the observation times, which are spaced at regular intervals over the duration of the simulation. The sequence of azimuths is used to draw a corresponding sequence of range-profiles from the appropriate data set in the URISD. The observations $r_k(t)$ are formed as the sum of these range-profiles and a sequence of complex, Gaussian distributed pseudo-random numbers with mean 0 and variance N_0 . The noisy observations are read from a file by our algorithm at the start of the simulation.

5.3.2 Libraries

A database is developed for each target that characterizes the set of range-profiles for that target at a discrete sampling of the possible azimuth angles. Each database is referred to as a library. When using the deterministic model, the libraries are created by selecting a fraction, typically one-half or one-fourth, of the range-profiles from the dataset from the URISD for each target at a fixed elevation angle and frequency band. When using the conditionally Gaussian model, the set of possible azimuth angles is divided into overlapping patches, and the libraries consist of mean range-profiles and range-bin variances computed from all of the range-profiles in each patch. The patch-size is arbitrarily chosen to be three degrees for the L-band data and one degree for the X-band data. Each patch overlaps with its neighbors by one-half of the patch size. The estimation algorithm searches through the libraries, seeking a sequence of library entries that maximize the joint likelihood for a particular sequence of observations. When using the conditionally Gaussian model, if the likelihood of observing a given range-profile is maximized by a particular azimuth patch, then the azimuth estimate

is set equal to the azimuth from the center of the patch. The nature of our search algorithm is detailed below.

5.3.3 Estimation Algorithm

The estimation algorithm we have developed for these simulations is based largely on our assumption that the azimuth varies linearly over a time window of length W , where W is an integer number of illuminations of the target. This constraint implies that a sequence of W azimuth estimates is determined by the estimates at the endpoints of the window. The estimation algorithm is initiated by conducting a series of global searches through the libraries for each candidate target to determine the vehicle and the azimuths at the window endpoints that maximize the joint likelihood for the first W observed range-profiles.

Future azimuth estimates are generated through a local search. The window endpoints are shifted by one illumination, and initial azimuth estimates at the new endpoints are found by linear interpolation from the previous estimates. Each of the endpoint estimates is constrained to be one of the discrete azimuths for which our libraries contain entries, and each estimate has two nearest neighbors in the libraries. Consider the current estimate and its neighbors as candidate azimuths, so that we have three candidates at each endpoint or nine possible pairs of candidates, each defining a linear orientation sequence over the window. The joint likelihood of each sequence is evaluated, using the current vehicle estimate, and the maximum likelihood candidate pair is selected as the new estimate. This process is repeated until the new estimate is equal to the current estimate. In this manner, the algorithm will proceed from the initial estimates produced by linear interpolation to the nearest local maximum of the joint likelihood surface.

Recognition is performed during the initialization process described above, in the sense that a global search is performed through the libraries for all four vehicles. To allow for the possibility that this search may have produced the incorrect vehicle estimate, the algorithm considers the hypothesis that an alternate vehicle is correct. The joint likelihood of the current azimuth estimates is computed for each alternate vehicle. The local search is again implemented to find the nearest local maximum of the likelihood surface. The joint likelihood under each hypothesis is computed, and the most likely is selected as the new estimate. This process is invoked at regular intervals defined by an integer number of illuminations. Given that our algorithm




Figure 5.6: Data from a typical realization of simulation 3. The left panel shows a sequence of 20 noise-corrupted L-band range profiles for the “m1” tank model. The signal to noise ratio is approximately 22 dB. The right panel shows a sequence of 20 library entries corresponding to our sequence of azimuth estimates at the observation times.

may produce different estimates of target type over the course of a single simulation, we define correct recognition as the case where the algorithm selects the correct vehicle during the initialization process, and maintains the correct vehicle estimate throughout the simulation.

5.3.4 Simulation Results

We have implemented our algorithm for a particular encounter scenario, using both the L-band and the X-band data from the URISD, and for a range of values of the noise variance N_0 . The scenario parameters were as follows. The target is the “m1” tank model. The target begins the simulation 500m west and 1000m north of the sensor, with an initial azimuth angle of 220° relative to the sensor, and moves in a straight line at 17.78m/s. The sensor moves north at 88.9m/s maintaining a constant elevation angle of 10° between itself and the target. The sensor illuminates the target 200 times at 50ms intervals.

Data from a typical realization of this scenario are shown in Figure 5.6. The left panel shows a sequence of 20 noise-corrupted L-band range profiles for the “m1” tank model. The right panel shows a sequence of 20 library entries corresponding to our sequence of azimuth estimates at the observation times. For this example, the noise variance is set to a value of 0.6 square meters. Let the signal to noise ratio be defined as the ratio of total signal energy to noise power. Because the energy in the




Figure 5.7: Estimation results from a typical realization of simulation 3 using L-band data. The true sequence of azimuth angles assumed by the target appears as a smooth curve, while the sequence of estimates appears as a jagged approximation to this curve. Note that both the true vehicle type and estimated vehicle type are the “m1” tank model, indicating successful recognition.

signal may change drastically over the course of the simulation, the SNR varies as well. The SNR for the data depicted in Figure 5.6 is approximately 22 dB.

The performance of our algorithm under the deterministic model for this example is shown in Figure 5.7. The true sequence of azimuth angles assumed by the target appears as a smooth curve, while the sequence of estimates appears as a jagged approximation to this curve. The two traces are difficult to distinguish in the figure as a result of the high quality of the azimuth estimates generated in this case. Note that both the true vehicle type and estimated vehicle type are the “m1” tank model, indicating successful recognition. The performance shown here is typical for our algorithm when using L-band data at this noise level, in that the sequence of azimuth estimates tracks the true azimuth sequence very closely, and the vehicle is correctly recognized throughout the simulation.

Performance Analysis: L-band

A more extensive examination of the performance of our algorithm was conducted by repeating this simulation 500 times using independent realizations of the pseudo-random noise sequences. Performance measures were computed over the 500 runs in terms of the fraction of runs in which correct recognition was achieved, and the




Figure 5.8: Performance measures for simulation 3 using L-band data. Shown in the left panel is the Hilbert-Schmidt norm-squared on azimuth estimation error as a function of the noise variance. The center panel shows the fraction of simulation runs in which correct recognition was achieved. The right panel shows the angular estimation error for the case where the correct vehicle type is assumed known. In each panel, the solid line indicates the performance under the deterministic model, and the dashed line indicates the performance under the conditionally Gaussian model.

Hilbert-Schmidt norm of the error in the azimuth estimates. This process was repeated for noise variances between 0.6 and 2.4 square meters, under both the deterministic model and the conditionally Gaussian model.

As discussed in Chapter 4, a performance measure for estimators in which the parameter is a rotation that takes its value on $SO(2)$ or $SO(3)$ is given by the Hilbert-Schmidt norm. For azimuth estimation on the circle, the square of the Hilbert-Schmidt norm is equal to

$$4 - 4 \cos(\theta_d), \quad (5.3)$$

where θ_d is the difference between the true azimuth angle and the estimate. We compute the Hilbert-Schmidt norm-squared for each azimuth estimate, and this performance measure is averaged over the number of estimates generated in each simulation run and over the 500 runs at each noise level to generate an estimate of the expected estimation error.

The results of these simulations for L-band data are shown in Figure 5.8. In the left panel of the figure, the Hilbert-Schmidt norm-squared on azimuth estimation error is plotted as a function of the noise variance. The center panel shows the fraction of simulation runs in which correct recognition was achieved. The right panel shows the angular estimation error for the case where the correct vehicle type is assumed known. In each panel, the solid line indicates the performance under the deterministic model, and the dashed line indicates the performance under the conditionally Gaussian model.

The figure indicates generally good recognition performance, in that the algorithm achieved correct recognition better than 90 percent of the time for noise variances up to 1.4 square meters, and better than 70 percent correct recognition for all noise variances, under either model. Also, the Hilbert-Schmidt norm-squared on the angular estimation error is consistently less than 1.6. To interpret this result, consider that the Hilbert-Schmidt bound for the worst-case estimator has a value of 4. If we assume that the targets are physically asymmetric objects, then the worst-case estimator randomly selects the azimuth estimate with uniform probability density over the circle.

For this simulation, the performance under the conditionally Gaussian model is approximately equivalent to that under the deterministic model. Such performance is achieved by the Gaussian model while using considerably smaller libraries. For L-band data, the deterministic model libraries each contain 1200 complex-valued range-profiles, whereas the Gaussian model libraries contain 240 complex-valued mean range-profiles and real-valued range-bin variances. Thus the total size of the L-band libraries for the Gaussian model is smaller than those for the deterministic model by a factor of $3\frac{1}{3}$.

The plots of Figure 5.8 also possess features that indicate less than complete confidence in the results obtained. For example, the left panel shows that the azimuth estimation error was not a strictly increasing function of the noise variance as we would expect. Also, for small noise variances, the azimuth estimation error is actually larger for the case where the correct target type is assumed known. We attribute both of these counterintuitive results to the fact that the data shown in the figure were compiled using only 500 runs of our simulation.

It must be noted that the performance of our algorithm is critically dependent on the choice of the size of the sliding window over which the linear azimuth dynamical model is enforced. For optimal recognition performance in the presence of noise, longer windows are desired. For the deterministic model, the log-likelihood for a sequence of observations is given by the sum of the individual log-likelihoods. Therefore, if the azimuth estimates are accurate, the use of longer windows increases the log-likelihood margin between the correct target and all other targets. However, because the true azimuth sequence is not linear, the use of longer windows increases the azimuth estimation error over the length of the window. This establishes a fundamental performance trade-off for the choice of this parameter. The results presented




Figure 5.9: Performance measures for simulation 3 using X-band data. Shown in the left panel is the Hilbert-Schmidt norm-squared on azimuth estimation error as a function of the noise variance. The center panel shows the fraction of simulation runs in which correct recognition was achieved. Note that correct recognition is achieved only when the algorithm produces the correct vehicle estimate throughout the entire simulation. The right panel shows the angular estimation error for the case where the correct vehicle type is assumed known.

above were obtained after trial-and-error suggested a window length of 20 observations provided the best performance when using the L-band data.

Performance Analysis: X-band

The simulations described above were repeated for identical scenario parameters using the X-band data from the URISD. The simulations were run 500 times at noise variances between 0.6 and 12.4 square meters. The recognition rates and Hilbert-Schmidt norms on the azimuth estimation errors were computed as before. The results of these simulations are shown in Figure 5.9. It is clear from the figure that the performance of the algorithm under the deterministic model suffered greatly when using X-band data. The left panel indicates that the azimuth estimation error exceeded the Hilbert-Schmidt bound for the worst case estimator at all noise variances tested. The center panel shows that correct recognition was never achieved for more than ten out of the 500 runs at any noise variance. Recall that correct recognition is achieved only when the algorithm produces the correct vehicle estimate throughout the entire simulation. The right panel indicates good performance in terms of azimuth estimation error at all noise levels, when the correct target type was assumed known. The poor performance observed here can be understood by examining Figure 3.5. The width of the central peak in the log-likelihood function is less than 0.1 degrees, and outside of this peak the log-likelihood for an incorrect target is larger than that for the correct

target. Thus, if the algorithm makes even a small error in the azimuth estimate, it will change the target type estimate, resulting in incorrect recognition.

It is also clear from the figure that the performance using X-band data was very significantly improved under the conditionally Gaussian model. A recognition rate of better than 70 percent was achieved, and the Hilbert-Schmidt norm-squared on the azimuth estimation error was consistently less than 1, for noise variances up to 8 square meters. This improved performance was achieved at a considerable savings in memory requirements over the deterministic model. The total size of the X-band data libraries for the conditionally Gaussian model is smaller than those for the deterministic model by a factor of more than 16.

5.4 Conclusions

This chapter has presented a series of three simulations designed to illustrate the potential for joint tracking and recognition using simulated HRR range-profiles. While the most complete performance analysis was performed for the last of these simulations, in each case we achieved general success. The implication of these results is that joint tracking and recognition from simulated range-profiles is possible, given an appropriate model for the the sensor data in the form of a likelihood function, data libraries that characterize the range profiles expected from a given target at the possible orientations, and an appropriate model for the target motion. The third simulation suggests that the conditionally Gaussian model provides improved performance over the deterministic model, especially at high frequencies.

Future simulations testing recognition could build on these efforts in many important ways. Of primary interest is adaptation of the simulations to operate on real HRR data. This would require a coordinated effort to obtain range-profiles for which ground truth values for the target shape and reflectivity, and the target orientations at the illumination times, are known. We are currently working with personnel in the radar community to obtain data for such a test. Also required would be libraries of simulated data for use with the conditionally Gaussian model. This, in turn, requires a verification of the range-profiles produced by simulators such as XPATCH.

Other factors that will be important to include in future ATR simulations are

- extension to full polarimetric HRR data,

- estimation of the complete target orientation, and
- inclusion of models with articulating parts.

Finally, and most importantly, will be the incorporation of the HRR processing used in these simulations into systems for joint detection, tracking and recognition of multiple targets where high resolution radar is only one of the sensors available.

Chapter 6

Conclusions and Future Directions

The research comprising this thesis has considered the problem of automatic target recognition using data from a high resolution radar sensor in the form of range-profiles. Specific attention has been focused on HRR models and optimal processing for recognition, issues associated with estimating orientation from range-profile data, and simulation experiments demonstrating joint tracking and recognition using HRR data.

6.1 Conclusions

In Chapter 1 the ATR problem was introduced. The general form of the problem involves multiple non-cooperative targets in a scene observed by multiple sensors of various types. The solution proposed in [31, 47, 76] was described, and it was seen that this algorithm is dependent on kinematic priors for target motion and statistical models for the observed data provided by the sensors. A simplified ATR problem was considered in which a single target is observed by a HRR sensor, and the goal is to jointly estimate its orientation and target type as it moves within the scene.

Chapter 2 provided an extensive review of the open literature pertaining to recognition using range-profile data. A variety of techniques for ATR were described, and several models for HRR data presented for use in ATR were examined. The general result of the literature search is that, while many investigators recognize the difficulty introduced to this problem by the well-known variability of range-profile data with target orientation, most ignore this problem by assuming that the orientation is known. Only Libby [45, 44] has considered the problem of jointly estimating orientation and target type from range-profile data, and his work does not present

results with respect to the accuracy of the orientation estimates. Additionally, most of the HRR models employed for this problem assume that the range-profile data can be described in terms of a low number of features that are derived intuitively and applied directly for recognition, instead of providing a direct statistical characterization of the observations.

The problem of statistical modeling of range-profiles was studied extensively in Chapter 3. A summary of basic concepts in high resolution radar was presented, including the WSSUS model introduced by Van Trees [86, 87]. A deterministic model for HRR data was discussed, and the likelihood function for varying target types and orientations was examined using simulated data. It was observed that by modeling the range-profile as deterministic, all of the variability of the data with orientation was represented in the likelihood function. Thus, while recognition with this model is seen to be possible, very accurate estimates of target orientation are necessary for recognition to be successful. A conditionally Gaussian model for the range-profile was introduced, and the likelihood function under this model was seen to be very promising for recognition, in that correct recognition is possible in the presence of significant errors in the orientation estimate. A series of recognition tests were conducted using this model, including tests that considered the use of full polarimetric data and variation in two Euler angles describing the target orientation. In every case where the two models were compared using these simple tests, it was seen that the recognition performance under the conditionally Gaussian model was more robust with respect to errors in the orientation estimate.

Chapter 4 examined a variety of issues relating to orientation estimation using HRR data. Global errors resulting from target symmetry and rotational ambiguity were described, and a technique for resolving the rotational ambiguity was investigated. The Hilbert-Schmidt bound on the error in orientation estimates was considered, and the bound was used to compare the performance under the two HRR models described in Chapter 3 and for two different HRR sensors. The results indicated better performance for the deterministic model in terms of orientation estimation, indicating the need for performance bounds for the joint estimation problem. A discussion of the importance of priors on the target kinematic state derived from dynamical models for the target motion was conducted. The efforts by Cutaia [23] and Srivastava [76] were cited as important contributions to the ATR problem. The priors they develop allow for elimination of unlikely target types through tracking measurements, and for

the constraining of likely future positions and orientations for the target given past values.

A series of three simulations was presented in Chapter 5. In the first of these, a single aircraft of known type was observed by a narrowband passive tracking array and a high resolution radar. The simulation demonstrated joint estimation of the target positions and orientations as it moved within the scene from the multiple sensor data. In the second simulation, a single aircraft was observed by an airborne sensor platform, and range-profile data were collected over a 10 second interval. The estimation algorithm jointly estimated the sequence of target orientations and distinguished between two candidate target types. The results for recognition for this simulation were very good, in that the algorithm consistently produced the correct target type estimate. The final simulation considered an air-to-ground scenario, in which an airborne sensor platform collects range profiles from a ground target, as both vehicles engage in straight-line motion. The algorithm produced estimates of the target orientations and discriminated among four candidate target types, using only the range-profile data. This analysis of this simulation was more comprehensive than that conducted for the previous two, in that the performance was studied as a function of the noise level and the frequency band of the simulated range-profiles. Also, the algorithm was tested using both of the HRR models developed in Chapter 3. The results of this final simulation indicated better performance under the conditionally Gaussian model, especially when using the X-band data. All three simulations demonstrated the potential for joint tracking and recognition, given an appropriate dynamical model for the target motion and an appropriate model for the HRR data in the form of a likelihood function.

Clearly there is much work to be done on this problem. We conclude this thesis by considering some specific directions for future work that have been revealed during the preparation.

6.2 Future Directions

The most important direction for future work on this problem is to apply the models developed in Chapter 3 to real range-profile data, and to adapt simulation like those in Chapter 5 to operate on these data. This prospect brings with it several challenges. The first of these is obtaining real data that would be appropriate for such analysis. Many of the high resolution radar systems from which such data might be obtained

are operated by the military, and as a result important parameters associated with the radar are often classified. Furthermore, obtaining range-profiles for which the ground truth in terms of the target orientations at the illumination times is exceedingly difficult to obtain. Nevertheless, we have made contacts in this regard with personnel at the U.S. Army Missile Command, and we are working with them toward obtaining data for such tests in the near future. Once the data are obtained, a primary issue will be how to construct libraries of range-profiles for use in estimation algorithms. If the libraries are to be constructed using real data, then the question arises as to what orientations are represented in the data, and how to address the problem of highly nonuniform sampling of the possible orientations. Also, if the ground truth orientations are known only to within, say, 5 degrees, then this poses an additional problem. If instead the libraries are to be constructed using data provided by a simulator such as XPATCH, then the question of the correspondence between real and simulated data must be addressed. We have no results showing the degree to which real and XPATCH simulated range-profiles for the same target and orientation compare.

In Chapter 3, we presented results indicating potentially better recognition performance when full polarimetric range-profiles are used, in conjunction with an appropriate structure for the polarimetric covariance. It would be worthwhile to study the improvement in performance resulting from inclusion of multiple polarizations through further simulations. This was not done in the simulations in Chapter 5 primarily out of a desire to keep the size of the range-profile libraries manageable.

The Hilbert-Schmidt bounds on errors in the estimation of orientation indicated better performance for the deterministic model than for the conditionally Gaussian model, for the single example that was considered. This result is in contrast to the analysis presented in Chapter 3, in which we concluded that the Gaussian model was more appropriate for recognition in the presence of errors in the orientation estimate through examination of several likelihood functions. To resolve this discrepancy, it would be very useful to obtain performance bounds on the probability of correct recognition achievable by any estimator using a given data set. Ideally, bounds on the performance of joint estimation of target type and orientation would be obtained, but we have completed no such analysis at this time.

An algorithm that is currently under study by the military for ATR using range-profiles is the Adaptive Gaussian Classifier or AGC. This algorithm is the subject of considerable attention in the classified literature. An outline of the salient

features of this algorithm is provided in Chapter 2, see also [26, 51]. This algorithm bears some resemblance to the algorithms described in this thesis, in that the Gaussian likelihood function is used for recognition. However, an important distinction is that a Gaussian model is applied to positive, real-valued data whose statistics are known to be non-Gaussian. Also, the AGC does not attempt to estimate the target orientation from the range-profile data, but rather requires that the orientation be provided externally. If the classified nature of this algorithm does not preclude it, there would be considerable interest in performing a direct comparison between the AGC and the algorithms presented in this thesis. Such a comparison would help to quantify the gain in performance achieved by jointly estimating orientation and target type, and would help to reveal if the increased algorithmic complexity is warranted.

Finally, and most importantly, it is desired to incorporate the range-profile processing demonstrated in this thesis into a system for joint detection, tracking, and recognition of multiple targets as outlined in [31, 47, 76]. Such a system represents a major advance in the state-of-the-art in ATR algorithms. As a result of the work of a multitude of researchers in the Center for Imaging Science, many of the pieces are already in place. The references cited above have clearly outlined the general ATR problem, the scene representation in terms of pattern theoretic constructs, and the solution in terms of jump-diffusion processes. Additional work has been conducted in the representation of forward-looking infrared or FLIR scenes, including: joint tracking and recognition of multiple ground vehicles from FLIR data; characterization of the variability in FLIR images in terms of an eigenbasis representation for the thermodynamic state of ground targets; and flexible shape modeling for rooftop recognition in FLIR scenes. These efforts and other work will be unified to produce a system for recognition of multiple targets, incorporating all of the variability present in natural scenes.

References

- [1] David C. Abell and Randy R. McElroy. Target orientation and the adaptive gaussian classifier. *Conversations*, November 1996.
- [2] Dennis J. Andersh, S. W. Lee, Robert Pinto, C. L. Yu, and Fredrick L. Beckner. Range profile synthesis for noncooperative target identification. Technical report, Wright Laboratories, 1991.
- [3] Dominick Andrisani, Frank P. Kuhl, and Daniel Gleason. A nonlinear tracker using attitude measurements. *IEEE Transactions on Aerospace and Electronic Systems*, 22(5):533-539, September 1986.
- [4] Kenneth Augustyn. A new approach to automatic recognition. *IEEE Transactions on Aerospace and Electronic Systems*, 28(1):105-114, January 1992.
- [5] John S. Baras and Sheldon I. Wolk. Model based automatic target recognition from high range resolution radar returns. *SPIE*, 2234:57-66, April 1994.
- [6] Jerold Baum, Eleanor Tung, and Steven Rak. Non-cooperative identification of ships with electrooptical data. *The Lincoln Laboratory Journal*, 7(1):3-29, 1994.
- [7] Peter N. Belhumeur, Joao P. Hespanha, and David J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. In *Proceedings of the European Conference on Computer Vision*, volume 1, pages 45-58, U.K., 1996.
- [8] Peter N. Belhumeur and David J. Kriegman. What is the set of images of an object under all possible lighting conditions? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1996.

- [9] Mark R. Bell and Robert K. Grubbs. JEM modeling and measurement for radar target identification. *IEEE Transactions on Aerospace and Electronic Systems*, 29(1):73-87, January 1993.
- [10] Bir Bhanu. Automatic target recognition: State of the art survey. *IEEE Transactions on Aerospace and Electronic Systems*, 22(4):364-379, July 1986.
- [11] Bir Bhanu and Terry L. Jones. Image understanding research for automatic target recognition. *IEEE Aerospace and Electronic Systems Magazine*, pages 15-22, October 1993.
- [12] Brett Borden. Problems in airborne radar target recognition. *Inverse Problems*, 10:1009-1022, 1994.
- [13] Maurice Borgeaud, Robert T. Shin, and Jin Au Kong. Theoretical models for polarimetric radar clutter. *Journal of Electromagnetic Waves and Applications*, 1(1):73-89, 1987.
- [14] W. M. Brown and C. W. Swonger. A prospectus for automatic target recognition. *IEEE Transactions on Aerospace and Electronic Systems*, 25(3):401-409, May 1989.
- [15] Barry D. Bullard and Patrick C. Dowdy. Pulse doppler signature of a rotary-wing aircraft. *IEEE Aerospace and Electronic Systems Magazine*, pages 28-30, May 1991.
- [16] Neil F. Chamberlain, Eric K. Walton, and Fred D. Garber. Radar target identification using polarization-diverse features. *IEEE Transactions on Aerospace and Electronic Systems*, 27(1):58-67, January 1991.
- [17] J. S. Chen and E. K. Walton. Comparison of two target classification techniques. *IEEE Transactions on Aerospace and Electronic Systems*, 22(1):15-21, January 1986.
- [18] Ken-Mu Chen, Dennis P. Nyquist, Edward J. Rothwell, Lance L. Webb, and Byron Drachman. Radar target discrimination by convolution of radar return with extinction-pulses and single mode extraction signals. *IEEE Transactions on Antennas and Propagation*, 34(7):896-904, July 1986.

- [19] Kun-Mu Chen, Dennis P. Nyquist, Edward J. Rothwell, and W. M. Sun. New progress on E/S pulse techniques for noncooperative target recognition. *IEEE Transactions on Antennas and Propagation*, 40(7):829-833, July 1992.
- [20] Dan S. Cheng and Allen R. Stubberud. Target tracking and identification using an extended-kalman-filter-based associative memory. In *Conference Record of the Twenty-Fifth Asilomar Conference on Circuits, Systems & Computers*, pages 604-607, San Jose, California, 1991. Maple Press, Inc.
- [21] Marvin N. Cohen. Variability of ultra-high range resolution radar profiles and some implications for target recognition. *SPIE*, 1699:256-266, 1992.
- [22] J. R. Copeland. Radar target classification by polarization properties. *Proceedings of the IRE*, 48(7):1290-1296, July 1960.
- [23] Nicholas J. Cutaia. *Performance of Automatic Target-Recognition Algorithms Using Kinematic Priors*. PhD thesis, Washington University, Saint Louis, Missouri, May 1996.
- [24] Mark R. Dewitt. High range resolution radar target identification using the Prony model and hidden markov models. Master's thesis, Air University, December 1992.
- [25] K. Cecilia Du. Range profile prediction in radar target identification. Master's thesis, Washington University, December 1993.
- [26] Christopher Lawrence Eisenbies. Classification of ultra high range resolution radar using decision boundary analysis. Master's thesis, Air University, December 1994.
- [27] Bernard Friedland. *Control System Design*, chapter 2. McGraw-Hill Book Company, New York, 1986.
- [28] Narendra S. Goel, editor. *Remote Sensing Reviews: Automatic Target Recognition and Tracking Based on Shape Analysis*, volume 6. Harwood Academic Publishers, 1992.
- [29] Thomas J. Green and Jeffery H. Shapiro. Maximum-likelihood laser radar range profiling with the expectation-maximization algorithm. *Optical Engineering*, 31(1):2343-2354, 1992.

- [30] U. Grenander. *General Pattern Theory*. Oxford University Press, 1993.
- [31] U. Grenander and M. I. Miller. Representations of knowledge in complex scenes. *Journal of the Royal Statistical Society B*, 56(3):549-603, 1994.
- [32] Gregory D. Hager and Peter N. Belhumeur. Real-time tracking of image regions with changes in geometry and illumination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1996.
- [33] Peter W. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 995-999, 1994.
- [34] Scott Hudson and Demetri Psaltis. Correlation filters for aircraft identification from radar range profiles. *IEEE Transactions on Aerospace and Electronic Systems*, 29(3):741-748, July 1993.
- [35] Malek G. M. Hussain. Principles of high-resolution radar based on nonsinusoidal waves: Part iii. radar-target reflectivity model. *IEEE Transactions on Electromagnetic Compatibility*, 32(2):144-152, May 1990.
- [36] Steven P. Jacobs, Joseph A. O'Sullivan, Mohammad Faisal, and Donald L. Snyder. Automatic target recognition system using high resolution radar. In *SPIE Proceedings*, volume 2845, Denver, Colorado, August 1996. The International Society for Optical Engineering.
- [37] Ismail Jouny and Fred D. Garber. M-ary sequential hypothesis tests for automatic target recognition. *IEEE Transactions on Aerospace and Electronic Systems*, 28(2):473-483, April 1992.
- [38] B. Kang, Bernard D. Steinberg, and S. B. Kesler. Comparison of 1-d and 2-d aircraft images. In *Proceedings: Unknown Conference*, pages 525-528, 1989.
- [39] J. D. Kendrick, P. S. Maybeck, and J. G. Reid. Estimation of aircraft target motion using orientation measurements. *IEEE Transactions on Aerospace and Electronic Systems*, 17(2):254-259, March 1981.
- [40] John D. Kraus and Keith R. Carver. *Electromagnetics*. McGraw-Hill, New York, second edition, 1973.

- [41] Aharon A. Ksienski, Yau-Tang Lin, and Lee James White. Low-frequency approach to target identification. *Proceedings of the IEEE*, 63(12):1651-1659, December 1975.
- [42] Frank P. Kuhl and R. Covelli. Object identification by multiple observations of the scattering matrix. *Proceedings of the IEEE*, 54:1110-1115, 1965.
- [43] Mun K. Leung and Thomas S. Huang. An integrated approach to 3-d motion analysis and object recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(10):1075-1084, October 1991.
- [44] Edmund W. Libby and Peter S. Maybeck. Application of sequence comparison techniques to multisensor data fusion and target recognition. In *Proceedings of the 32nd Conference on Decision and Control*, December 1993.
- [45] Edward W. Libby. *Application of Sequence Comparison Methods to Multisensor Data Fusion and Target Recognition*. PhD thesis, Air University, June 1993.
- [46] O. Lowenschuss. Scattering matrix application. *Proceedings of the IEEE*, 54:988-992, 1965.
- [47] M. I. Miller, U. Grenander, J. A. O'Sullivan, and D. L. Snyder. Automatic target recognition organized via jump-diffusion algorithms. *IEEE Transactions on Image Processing, Special Issue on Automatic Target Detection and Recognition*, 1997. Accepted for Publication September 1, 1996.
- [48] M. I. Miller, A. Srivastava, and U. Grenander. Conditional-mean estimation via jump-diffusion processes in multiple target tracking/recognition. *IEEE Transactions on Signal Processing*, 43(11):2678-2690, November 1995.
- [49] M.I. Miller and D. R. Fuhrmann. Maximum likelihood narrow-band direction finding and the EM algorithm. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 38(9):560-577, 1990.
- [50] Michael I. Miller, Robert S. Teichman, Anuj Srivastava, Joseph A. O'Sullivan, and Donald L. Snyder. Jump-diffusion processes for automated tracking-target recognition. In *1993 Conference on Information Sciences and Systems*, Baltimore, Maryland, March 1993. Johns Hopkins University.

- [51] Richard A. Mitchell and Rob DeWall. Overview of high range resolution radar target identification. Technical report, Wright Laboratories, 1994.
- [52] David L. Moffat and Chiwei Chuang. Natural resonances of radar targets via Prony's method and target discrimination. *IEEE Transactions on Aerospace and Electronics Systems*, 12(5):583-589, September 1976.
- [53] David L. Moffat and Richard K. Mains. Detection and discrimination of radar targets. *IEEE Transactions on Antennas and Propagation*, 23(3):358-367, May 1975.
- [54] David L. Moffat, Johnathan D. Young, Aharon A. Ksienski, Heng-Cheng Lin, and Charles M. Rhoads. Transient response characteristics in identification and imaging. *IEEE Transactions on Antennas and Propagation*, 29(2):192-205, March 1981.
- [55] Hiroshi Murase and Shree K. Nayar. Visual learning and recognition of 3-d objects from appearance. *International Journal of Computer Vision*, 14(1):5-24, 1995.
- [56] Fred E. Nathanson. *Radar Design Principles*. McGraw-Hill, New York, 1991.
- [57] Leslie M. Novak and Michael C. Burl. Optimal speckle reduction in polarimetric SAR imagery. *IEEE Transactions on Aerospace and Electronic Systems*, 26(2):293-305, March 1990.
- [58] Leslie M. Novak, Michael C. Burl, and W. W. Irving. Optimal polarimetric processing for enhanced target detection. *IEEE Transactions on Aerospace and Electronic Systems*, 29(1):234-243, January 1993.
- [59] Leslie M. Novak, Michael B. Sechtin, and Michele J. Cardullo. Studies of target detection algorithms that use polarimetric radar data. *IEEE Transactions on Aerospace and Electronic Systems*, 25(2):150-165, March 1989.
- [60] Joseph A. O'Sullivan, K. Cecilia Du, Robert S. Teichman, Michael I. Miller, Donald L. Snyder, and Vincent C. Vannicola. Radar target recognition using shape models. In *Proceedings of the 30th Annual Allerton Conference on Communication Control and Computing*, pages 515-523, Urbana, Illinois, October 1992. University of Illinois.

- [61] Joseph A. O'Sullivan, K. Cecilia Du, Robert S. Teichman, Michael I. Miller, Donald L. Snyder, and Vincent C. Vannicola. Reflectivity models for radar target recognition. In *SPIE Proceedings*, volume 1960, pages 152-161, Orlando, Florida, April 1993. The International Society for Optical Engineering.
- [62] Joseph A. O'Sullivan, Steven P. Jacobs, Michael I. Miller, and Donald L. Snyder. A likelihood-based approach to joint target tracking and identification. In Avtar Singh, editor, *Conference Record of the Twenty-Seventh Asilomar Conference on Signals, Systems & Computers*, pages 290-294, Los Alamitos, California, November 1993. IEEE Computer Society Press.
- [63] Alex Pentland, Baback Moghaddam, and Thad Starner. View-based and modular eigenspaces for face recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 84-91, 1994.
- [64] August W. Rihaczek. *Principles of High-Resolution Radar*. Peninsula Publishing, Los Altos, California, 1985.
- [65] Edward J. Rothwell, Kun-Mu Chen, and Dennis P. Nyquist. Extraction of the natural frequencies of a radar target from a measured response using E-pulse techniques. *IEEE Transactions on Antennas and Propagation*, 35(6):715-720, June 1987.
- [66] Edward J. Rothwell, Dennis P. Nyquist, Kun-Mu Chen, and Byron Drachman. Radar target discrimination using the extinction-pulse technique. *IEEE Transactions on Antennas and Propagation*, 33(9):929-936, September 1985.
- [67] R. Schmidt. *A signal subspace approach to multiple emitter location and spectral estimation*. PhD thesis, Stanford University, Palo Alto, CA, November 1981.
- [68] J. A. Shapiro, B. A. Capron, and R. C. Harney. Imaging and target detection with a heterodyne-reception optical radar. *Applied Optics*, 20(19):3292-3313, October 1981.
- [69] Jeffery A. Shapiro. Target-reflectivity theory for coherent laser radars. *Applied Optics*, 21(18):3398-3407, September 1982.
- [70] Jeffery H. Shapiro. Correlation scales of laser speckle in heterodyne detection. *Applied Optics*, 24(12):1883-1888, June 1985.

- [71] Paul B. Silverstein, O. Scott Sands, and Fred D. Garber. Radar target classification and interpretation by means of structural descriptions of backscatter signals. *IEEE Aerospace and Electronic Systems Magazine*, pages 3-7, May 1991.
- [72] Merrill I. Skolnik. *Introduction to Radar Systems*, chapter 2. McGraw-Hill Book Company, second edition, 1980.
- [73] C. Ray Smith and Paul M. Goggans. Radar target identification. *IEEE Antennas and Propagation Magazine*, 35(2):27-37, April 1993.
- [74] Donald L. Snyder, Joseph A. O'Sullivan, and Michael I. Miller. The use of maximum likelihood estimation for forming images of diffuse targets from delay-doppler data. *IEEE Transactions on Information Theory*, 35(3):536-548, May 1989.
- [75] Anuj Srivastava. Automated target tracking and recognition using jump-diffusion processes. Master's thesis, Washington University, December 1993.
- [76] Anuj Srivastava. *Inferences on Transformation Groups Generating Patterns on Rigid Motions*. PhD thesis, Washington University, Saint Louis, Missouri, July 1996.
- [77] Anuj Srivastava, Nicholas J. Cutaia, Michael I. Miller, Joseph A. O'Sullivan, and Donald L. Snyder. Multi-target narrowband direction finding and tracking using motion dynamics. In *Proceedings of the 30th Annual Allerton Conference on Communication Control and Computing*, pages 279-288, Urbana, Illinois, October 1992. University of Illinois.
- [78] Anuj Srivastava, Michael I. Miller, and Ulf Grenander. Multiple target direction of arrival tracking. *IEEE Transactions on Signal Processing*, 43(5):1282-85, May 1995.
- [79] William M. Steedly and Randolph L. Moses. High resolution exponential modeling of fully polarized radar returns. *IEEE Transactions on Antennas and Propagation*, 27(3):459-468, May 1991.
- [80] Bernard D. Steinberg. *Principles of Aperture and Array System Design*, chapter 8. John Wiley & Sons, 1976.

- [81] Bernard D. Steinberg. Target recognition and detection sensitivity in two-dimensional microwave imaging. Technical Report UP-VFRC-2-90, RADC/OCD Griffiss Air Force Base, Rome, New York, January 1990.
- [82] D. D. Sworder, P. F. Singer, D. Doria, and R. G. Hutchins. Image-enhanced estimation methods. *Proceedings of the IEEE*, 81(6):797-812, June 1993.
- [83] Robert S. Teichman. Automated target recognition in a distributed computing environment. Master's thesis, Washington University, September 1994.
- [84] Charles W. Therrien. A sequential approach to target discrimination. *IEEE Transactions on Aerospace and Electronic Systems*, 14(3):433-440, May 1978.
- [85] R. Touzi and A. Lopes. The principle of speckle filtering in polarimetric SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 32(5):1110-1114, September 1994.
- [86] Harry L. Van Trees. *Detection, Estimation and Modulation Theory*, volume 1. John Wiley and Sons, New York, 1968.
- [87] Harry L. Van Trees. *Detection, Estimation and Modulation Theory*, volume 3. John Wiley and Sons, New York, 1971.
- [88] Donald R. Wehner. *High Resolution Radar*. Artech House, Norwood, Massachusetts, 1987.
- [89] L. J. White and A. A. Ksienski. Aircraft identification using a bilinear surface representation of radar data. *Pattern Recognition*, 6:35-45, 1974.
- [90] Anthony Zyweck and Robert E. Bogner. Radar target recognition using range profiles. In *Proceedings of the 1994 International Conference on Acoustics, Speech & Signal Processing*, volume 2, pages 373-376. Institute of Electrical and Electronics Engineers, Signal Processing Society, April 1994.
- [91] Anthony Zyweck and Robert E. Bogner. Coherent averaging of range profiles. In *Record of the IEEE 1995 International Radar Conference*, pages 456-61, 794, Alexandria, Virginia, May 1995. Institute of Electrical and Electronics Engineers.
- [92] Anthony Zyweck and Robert E. Bogner. Radar target classification of commercial aircraft. *IEEE Transactions on Aerospace and Electronic Systems*, 32(2):598-606, April 1996.

Vita

Steven P. Jacobs

Born

18 November 1966, Arlington Heights, IL, USA.

Education

1993-96:

Doctor of Science in Electrical Engineering

Washington University, St. Louis, Missouri

Dissertation: Automatic Target Recognition Using Sequences of High Resolution Radar Range-Profiles

1990-93:

Northwestern University, Evanston, Illinois

Doctoral study in Systems for Distributed Detection and Estimation

Transferred to Washington University, August 1993.

1988-90:

Master of Science in Electrical Engineering

Washington University, St. Louis, Missouri

Thesis: Parallel Implementation of an EM Algorithm for a Radar Imaging Problem

1984-87:

Bachelor of Science in Electrical Engineering

Washington University, St. Louis, Missouri

Areas of Interest

Estimation and communications theory, radar systems, signal processing, automatic target recognition

Conference Publications

S. P. Jacobs, J. A. O'Sullivan, and Z. Bekker, "Joint Tracking and Recognition of Ground Targets from High Resolution Radar Range-Profiles," To Appear, *Proceedings of the Third Workshop on Conventional Weapon ATR*, Huntsville, Alabama, November, 1996.

S. P. Jacobs, J. A. O'Sullivan, M. Faisal, and D. L. Snyder, "Automatic Target Recognition System Using High Resolution Radar," *Proceedings of the 1996 SPIE*

International Symposium on Optical Science, Engineering, and Instrumentation, Volume 2845: 323-336, August 1996.

J. A. O'Sullivan, S. P. Jacobs, M. I. Miller, and D. L. Snyder, "A Likelihood-Based Approach to Joint Target Tracking and Identification," *Conference Record: Twenty-Seventh Asilomar Conference on Signals, Systems & Computers*, 290-294, November 1993.

J.A. O'Sullivan, P. Moulin, D. L. Snyder, and S. P. Jacobs, "Computational Considerations for Maximum-Likelihood Radar Imaging," *Proceedings of the 1990 Conference on Information Science and Systems*, March 1990.

Other Publications

S. P. Jacobs, J. A. O'Sullivan, M. Faisal, A. Srivastava, D. L. Snyder, and M. I. Miller, "Automatic Target Recognition Using Sequences of High Resolution Radar Range-Profiles," *ESSRL Monograph*, Number ESSRL-95-26, November 1995.

Teaching Experience

- 1996: Instructor, undergraduate course:
"Signal Analysis for Electronic Systems and Circuits"
- 1996: Guest Lecturer, two lectures in graduate course:
"Radar Systems"
- 1996: Guest Lecturer, one lecture in graduate course:
"Detection and Estimation Theory"
- 1995: Guest Lecturer, nine lectures in undergraduate course:
"Signal Analysis for Electronic Systems and Circuits"
- 1994: Guest Lecturer, two lectures in graduate course:
"Probability and Stochastic Processes"
- 1991-1993: Laboratory Instructor, undergraduate course:
"Linear Circuit Theory"
- 1990: Teaching Assistant, grading and consulting for graduate course:
"Probability and Random Variables"
- 1989: Teaching Assistant, grading and consulting for undergraduate course:
"Communication Theory"
- 1988: Instructor (Watterson College), two technical courses:
"Linear Circuit Theory" and "Introduction to Robotics"

Affiliations

Institute of Electrical and Electronics Engineers
Aerospace and Electronic Systems Society; Signal Processing Society

Eta Kappa Nu
Tau Beta Pi

Scholarships and Fellowships

1990-91: Murphy Fellowship, Northwestern University
1989-90: Mudd Fellowship, Washington University

May 1997

Short Title: ATR Using HRR Range-Profiles

Jacobs, D.Sc. 1997